



REMEDIES AGAINST

THE VOCABULARY GAP

IN INFORMATION RETRIEVAL

**CHRISTOPHE
VAN GYSEL**

Remedies against the Vocabulary Gap in Information Retrieval

Christophe Van Gysel

This version differs from the official version submitted to the university
through minor changes in Chapters 4 and 8.

Remedies against the Vocabulary Gap in Information Retrieval

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex
ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in
de Agnietenkapel
op vrijdag 17 november 2017, te 12:00 uur

door

Christophe Jacky Henri Van Gysel

geboren te Lier, België

Promotiecommissie

Promotor:

Prof. dr. M. de Rijke Universiteit van Amsterdam

Co-promotor:

Dr. E. Kanoulas Universiteit van Amsterdam

Overige leden:

Prof. dr. B. Goethals Universiteit Antwerpen

Dr. K. Hofmann Microsoft Research Cambridge

Dr. C. Monz Universiteit van Amsterdam

Prof. dr. M. Welling Universiteit van Amsterdam

Prof. dr. M. Worring Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

The research was supported by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement nr 312827 (VOX-Pol), the Google Faculty Research Award scheme and the Bloomberg Research Grant program.

All content represents the opinion of the author, which is not necessarily shared or endorsed by his respective employers and/or sponsors.

Computing resources were provided by the Netherlands Organisation for Scientific Research (NWO) through allocation SH-322-15 of the Cartesius system, the Advanced School for Computing and Imaging (ASCI) by allocation of the Distributed ASCII Supercomputer 4 (DAS-4) system and the Information and Language Processing Systems group.

Copyright © 2017 Christophe Van Gysel, Amsterdam, The Netherlands

Cover by Samira Abnar and Mostafa Dehghani

Printed by Off Page, Amsterdam

ISBN: 978-94-6182-838-5

Acknowledgements

I never intended to pursue a doctorate degree.

Senior researchers—academics and industry folks alike—kept repeating that a masters degree and an engineering position at a high-end technology giant would be the better career path. After all, postgraduate studies are often described as one of the few remaining forms of modern slavery in the Western world.¹ Given that this dissertation completes my postgraduate degree, I can only conclude that I'm susceptible to reverse psychology.

At the end of August 2014, I was spending a few months in New York. It was there that I was told that the postgraduate student experience has just one significant factor: the student's advisor. This was only a few days after I had first met with my prospective advisor—Maarten. The postgraduate experience boiled down to the following question: *"Is the student's advisor a lion or a wolf?"*² Will your advisor fight for you—as a lion—or misguide you and stab you in the back—as a wolf? After these three years, I can say with certainty that Maarten resembles something closer to a griffin: a mythical lion with wings. Consequently, it goes without saying that I express my deepest gratitude towards Maarten for his guidance and ideas.

I also thank my co-advisor, Evangelos, for his insightful feedback and the inspiring conversations we shared. Before Evangelos became my co-advisor, I was working closely with Marcel. Marcel, thank you for the guidance during those first few months.

I am honoured to have a graduation committee consisting of very talented researchers from a wide variety of backgrounds. Bart, Christof, Katja, Marcel, Max, thank you for taking the time to read my dissertation and your valuable feedback. Likewise, I thank my paronyms, Hosein and Rolf, for standing by me during the defence of this dissertation.

I thank the vocabulary gap for being an interesting problem that I could write this dissertation on.

Science cannot exist without collaborations and I thank my co-authors for all the hard work. Alexey, Bart, Bhaskar, Evangelos, Françoise, Grzegorz, Ian, Leonid, Maarten, Marcel, Matteo, Mostafa, Nicola, Piotr, Roy and Tom, thank you for discussions, the modelling, the experiments, the writing, the polishing and the publishing.

We all have to start somewhere and I would express my gratitude to the people who helped me during the early days of my post-graduate studies. Thank you, Daan, Manos, Tom and Zhaochun.

Science, and consequently life, would be boring without the social aspect of it. Thanks to my colleagues at ILPS for the insightful discussions, the support and the fun evenings: Abdo, Adith, Aldo, Aleksandr, Alexey, Ana, Anna, Anne, Arianna, Artem, Bob, Boris, Chang, Chuan, Cristina, Daan, Damien, Dan, Dat, David, David, Dilek, Eva, Evgeny, Fei, Hamid, Harrie, Hendrik, Hendrike, Hosein, Ilya, Isaac, Iva, Ivan, Julia, Kaspar,

¹<http://www.independent.co.uk/news/education/education-news/postgraduate-students-are-being-used-as-slave-labour-7791509.html>

²<https://www.math.ku.edu/~jmartin/fun/grad.html>

Katya, Ke, Lars, Maarten, Marlies, Marzieh, Masrour, Mostafa, Nikos, Praveen, Richard, Ridho, Rolf, Shangsong, Svitlana, Tobias, Tom, Xiaojuan, Xinyi, Yaser, Zhaochun and Ziming.

Petra, thank you for everything. You are invaluable within ILPS.

During the course of my masters studies at the University of Antwerp and my doctoral studies at the University of Amsterdam I gained priceless experience as part of my 7 internships: (1 & 2) Gus, Sam and Sam, with whom I worked at the ads ranking quality team of Google, thank you. It was during my time there that I first started wondering about pursuing a doctoral degree. (3) Alec, Chad and Christopher, who supervised me at the infrastructure security team at Facebook. Thanks for the hacking. (4) Thanks to Leonid, Ian and Françoise who provided me with my first industrial research experience at the speech team of Google. (5) Yi, Xiaochuan and Ilya, who motivated me to push the limits of speech recognition, at the language modelling team of Apple. (6) Bhaskar, Matteo and Nicola, whom I worked with at the query formulation team of Microsoft Bing. Thanks for including me into the team and showing me that research and product development go hand in hand. (7) Finally, thanks to Sushobhan, Xinran and Jie, at the search team of Snap Inc., who had a *vocabulary gap* that needed to be bridged.

Many friends supported my throughout my doctoral studies. I thank my friends that I met during my journey at the University of Antwerp and the Association of Mathematics, Informatics and Physics. It has been an amazing few years and I know that I should visit more often. Beyond Antwerp, I thank my friends that I met while travelling and participating in foreign exchanges.

I thank my parents, Catherine and Marc, for their eternal support. Without their guidance, I would not be where I am today. My brother, Cedric, thank you for your support and the discussions. My grandparents, Harry, Henri, Jacky, Lisette, thank you for always looking after me. My aunt and uncle, Ingrid and Patrick, thank you for conversations and gatherings. My cousins, Alexandar, Amelie, Laurens, Sarah, thanks for all the good times.

Last, but not least, I thank my girlfriend, Katya, for all her love, support and understanding.

1	Introduction	1
1.1	Research outline and questions	2
1.1.1	Query formulation	2
1.1.2	Latent vector spaces	3
1.2	Main contributions	4
1.2.1	Algorithmic contributions	4
1.2.2	Theoretical contributions	4
1.2.3	Empirical contributions	5
1.2.4	Open-source software contributions	5
1.3	Thesis overview	6
1.4	Origins	6
2	Background	9
2.1	Query formulation	9
2.1.1	Session search	9
2.1.2	Proactive information retrieval	10
2.1.3	Predictive models for email	10
2.1.4	Query formulation and reformulation	11
2.2	Latent semantic models	12
2.2.1	Latent semantic models for information retrieval	12
2.2.2	Entity retrieval	13
2.2.3	Neural language modelling	15
2.2.4	Neural information retrieval	16
I	Query Formulation	17
3	Lexical Query Modelling in Session Search	19
3.1	Introduction	19
3.2	Lexical matching for sessions	19
3.3	Experiments	22
3.3.1	Benchmarks	22
3.3.2	Evaluation measures	22
3.3.3	Systems under comparison	22
3.3.4	Ideal lexical term weighting	23
3.4	Results & discussion	23
3.5	Summary	27
4	Reply With: Proactive Recommendation of Email Attachments	29
4.1	Introduction	29
4.2	Related work	31
4.3	Proactive attachable item recommendation	32
4.3.1	Attachment retrieval	33
4.3.2	Evaluating query formulations	34

4.4	Query formulation model	34
4.4.1	Model training	35
4.4.2	A convolutional neural network for ranking query terms . . .	37
4.5	Experimental set-up	40
4.5.1	Research questions	40
4.5.2	Experimental design	41
4.5.3	Data collections and pre-processing	41
4.5.4	Methods under comparison	42
4.5.5	Evaluation measures and significance	43
4.6	Results & discussion	43
4.6.1	Overview of experimental results	43
4.6.2	Analysis of differences	46
4.6.3	Feature importance	47
4.7	Summary	48

II Latent Vector Spaces 49

5	Unsupervised, Efficient and Semantic Expertise Retrieval	51
5.1	Introduction	51
5.2	Related work	52
5.3	A log-linear model for expert search	53
5.3.1	The model	53
5.3.2	Parameter estimation	54
5.4	Experimental setup	55
5.4.1	Research questions	55
5.4.2	Benchmarks	56
5.4.3	Baselines	56
5.4.4	Implementation details	58
5.5	Results & discussion	58
5.5.1	Overview of experimental results	58
5.5.2	Error analysis	63
5.5.3	Scalability and efficiency	65
5.5.4	Incremental indexing	69
5.6	Summary	69
5.7	Appendix	70
6	Structural Regularities in Text-based Entity Vector Spaces	73
6.1	Introduction	73
6.2	Related work	74
6.3	Text-based entity vector spaces	75
6.4	Experimental set-up	76
6.4.1	Research questions	76
6.4.2	Expert finding collections	77
6.4.3	Implementations and parameters	77
6.5	Regularities in entity vector spaces	78

6.5.1	Answers to research questions	78
6.5.2	Analysis of the expert prior in the log-linear model	82
6.6	Summary	84
7	Learning Latent Vector Spaces for Product Search	87
7.1	Introduction	87
7.2	Related work	89
7.3	Latent vector spaces for entity retrieval	90
7.3.1	Background	90
7.3.2	Latent semantic entities	91
7.3.3	Parameter estimation	92
7.4	Experimental setup	93
7.4.1	Research questions	93
7.4.2	Experimental design	94
7.4.3	Product search benchmarks	97
7.4.4	Evaluation measures and significance	97
7.4.5	Methods used in comparisons	97
7.5	Results & discussion	98
7.5.1	Overview of experimental results	98
7.5.2	A feature for machine-learned ranking	100
7.6	Analysis of representations	102
7.7	Summary	105
8	Neural Vector Spaces for Unsupervised Information Retrieval	107
8.1	Introduction	107
8.2	Related work	108
8.3	Learning semantic spaces	109
8.3.1	The Neural Vector Space Model	109
8.3.2	The objective and its optimization	110
8.3.3	Implementation	113
8.4	Experimental setup	114
8.4.1	Research questions	114
8.4.2	Benchmark datasets & experiments	115
8.4.3	Retrieval models considered for comparison	116
8.4.4	Evaluation measures and statistical significance	119
8.5	Results	120
8.5.1	Performance of NVSM	120
8.5.2	Query-level analysis	120
8.5.3	Semantic vs. lexical matching	120
8.5.4	NVSM and Luhn significance	125
8.6	Discussion and analysis	125
8.6.1	An investigation of judgement bias	125
8.6.2	Unsupervised deployment	126
8.7	Summary	128

9	Conclusions	135
9.1	Main findings	135
9.2	Future work	137
A	Pyndri: A Python Interface to the Indri Search Engine	143
A.1	Introduction	143
A.2	Introducing Pyndri	143
A.2.1	Low-level access to document repository	144
A.2.2	Querying Indri from Python	145
A.3	Summary	145
B	Semantic Entity Retrieval Toolkit	147
B.1	Introduction	147
B.2	The toolkit	147
B.2.1	Collection parsing and preparation	148
B.2.2	Representation learning	148
B.2.3	Entity ranking & other uses of the representations	149
B.3	Summary	150
	Bibliography	151
	Samenvatting	162
	Summary	164

1

Introduction

Search engines heavily rely on term-based approaches according to which queries and documents are represented as bags of words. Text—a document or a query—is represented by a bag of its words that ignores grammar and word order, but retains word frequency counts. When presented with a search query, the engine then ranks documents according to their relevance scores by computing, among other things, the matching degrees between query and document terms. While term-based approaches are intuitive and effective in practice, they are based on the hypothesis that documents that exactly contain the query terms are highly relevant regardless of query semantics. Inversely, term-based approaches assume documents that do not contain query terms to be irrelevant. Li and Xu [146] note that a high matching degree at the term level does not necessarily mean high relevance and, vice versa, documents that match null query terms may still be relevant. Consequently, there exists a *vocabulary gap* between queries and documents that occurs when both use different words to describe the same concepts.

In addition to the crude heuristic that mandates document relevance to be a function of query/document term overlap, the ubiquity of term-based approaches can be explained by efficiency constraints imposed on retrieval engines. The typical concise nature of queries, together with term-based matching, is used to filter out documents without query terms using a specialized data structure called the *inverted index* [66, p. 129: Section 5.3]. An inverted index operates in a similar way as a subject index in a reference book. For example, in an encyclopedia, the subject index contains references to pages that discuss a particular subject. Consequently, the reader is redeemed of the cumbersome task of determining the relevance of every page individually w.r.t. her information need. Search engines employ a filter-and-refine strategy where an initial method, typically the lookup in an inverted index, is used to generate a pool of candidate documents [151, p. 135: Section 10.3]. The pool of candidate documents is subsequently re-ranked using more expensive methods [166].

There are, however, two major drawbacks of the inverted index that are relevant to this dissertation: (1) Following up on our earlier discussion, term-based retrieval may incorrectly classify relevant documents that do not contain query terms as irrelevant. This negatively affects the recall of the retrieval results. Given that term-based retrieval is used for an initial filtering of documents, semantic matching re-rankers applied during the refining step of the retrieval pipeline are powerless when it comes to relevant

documents whose terms overlap little with the textual query issued by the user. (2) To alleviate the filtering of false negatives, search engines often automatically reformulate queries by adding terms (i.e., expansion) [66, p. 199: Section 6.2.3]. However, query expansion is a double-edged sword as Grossman et al. [98] show that the number of CPU operations and the disk I/O grow superlinearly—while the number of novel discovered relevant documents diminishes rapidly—with the query size. In fact, incautious and overzealous query expansion causes query drift [180] and generates a candidate document set that is unrelated to the original query [165, p. 187: Section 9.1.6].

This dissertation directly targets these two pitfalls of the inverted index and approaches the shortcomings of the inverted index from two opposite angles. (1) In Part I, we formulate queries from complex and heterogeneous textual structures (search sessions and email threads) so as to fulfill the user’s information need. Search sessions consist of a sequence of user interactions (i.e., query reformulations, clicks), search engine responses (i.e., ranked documents in a result page, also referred to as SERP) and are indicative of a complex information need. The task is then to formulate a textual query to satisfy the overall information need that characterizes the session. In the case of email threads, the explicit information need is less straightforward. We focus on a particular case where an incoming email message is a request for content. The task is then to formulate a query that correctly retrieves the appropriate item from a document repository. Overall, the common goal of the research performed in Part I is to formulate queries that are (a) shorter and (b) exhibit more effective retrieval than taking the full textual structure as a query (i.e., queries/SERPs and full email messages for the session and email domains, respectively). (2) Query expansion is a slippery slope as adding more terms increases latency and can cause query drift. However, nearest neighbour algorithms can be more efficient than classical term-based retrieval [43] and, consequently, may be used as an alternative [147] or in addition [201] to the inverted index. Therefore, Part II is dedicated to the modelling of latent vector spaces for information retrieval (IR). In particular, we focus on retrieval domains where semantic matching is known to be important: entities [20, 69] and news articles [107].

1.1 Research outline and questions

The broad theme of this dissertation involves fulfilling information needs that are (1) embedded within complex textual structures, such as sessions or email threads through **query formulation**, and (2) expressed as textual queries that require **semantic matching** (e.g., informational queries [44]). Below, we introduce the high-level research questions answered in the respective chapters. In each chapter, we pose multiple sub-questions whose answers are combined to answer the higher-level questions below.

1.1.1 Query formulation

Information needs are often posed in a form that is more complex than a single short user-formulated textual query. What if we know the user is trying to accomplish a complex task by issuing multiple short queries? At first, the user issues an initial query. Subsequently, at each step the user observes feedback from the search engine

and reformulates her request. How can we use the previously-issued queries and the documents observed by the user to improve retrieval effectiveness? In the case of email threads, many emails contain implicit or explicit requests for content. Can we formulate a query from email threads that retrieves the relevant content to be attached? This brings us to our first research question:

RQ1 How to formulate a query from complex textual structures—such as search sessions or email threads—in order to better answer an information need?

We perform an analysis on session search logs from TREC and provide empirical insight in the potential of lexical language models (Chapter 3). Building upon this insight, we propose a Convolutional Neural Network (CNN) that formulates a query from email requests for the task of proactive attachment recommendation (Chapter 4).

1.1.2 Latent vector spaces

Entities (e.g., people, products) are often characterized by large bodies of text [23, 170]. Bag of words approaches may degrade for long documents due to their verbosity and scope [67]. In addition, entity domains require more semantic matching than domains where navigational queries are prevalent (e.g., Web search) as entity-oriented queries often describe the entity rather than searching for a known-item. For example, in expert finding, users describe the expertise of the expert they are searching for instead of the name of the expert herself [245]. When searching for products on an e-commerce website, users often formulate queries by listing the characteristics of the product they are interested in [217].

RQ2 Can we learn a latent vector space of retrievable entities that performs well for retrieval?

We introduce the log-linear model for expert finding and show its effectiveness on three expert finding benchmarks (Chapter 5). In expert finding, the user issues a topical query and is presented with a ranking of people. We perform an analysis of the semantic matching performed by the model and give insight in the regularities contained within latent expert representations (Chapter 6). For example, we show that experts who operate in similar domains have similar representations.

In **RQ2** we consider a particular entity domain (i.e., expert finding) that consists of a small number of entities that each have a large body of associated text. In addition, the queries in the expert domain require semantic matching, as users describe the domain of the expert and known-item queries are inherently not part of the expert finding task. However, the training procedure proposed in response to **RQ2** is linear in the number of entities. This is impractical as it severely limits the training speed. In the next research question, we address this impracticality and investigate sampling methods to scale up training to large entity spaces.

RQ3 Can we scale up latent vector spaces to larger entity domains that have less textual content per entity compared to the expert finding setting?

To answer **RQ3** we scale up the model training to large entity spaces by sampling negative examples (Chapter 7). However, its retrieval effectiveness diminishes as the number of retrievable entities increases. Furthermore, how do our latent vector spaces perform on non-entity domains? We address these concerns in **RQ4**.

RQ4 Can we further increase the retrieval effectiveness and scale latent vector spaces up to hundreds of thousands of documents?

Our final research question, **RQ4**, is answered in Chapter 8 by the introduction of the Neural Vector Space Model (NVSM). We improve the loss function of our latent vector spaces by incorporating IR-specific regularities and evaluate NVSM on article retrieval benchmarks from TREC.

1.2 Main contributions

The main contributions of this dissertation are listed in this section. Our contributions come in the form of algorithmic, theoretical, empirical and open-source software contributions. For each contribution, we list the chapter where the contribution is made or, in the case of software packages, where the package was used to generate experimental results.

1.2.1 Algorithmic contributions

1. A frequency-based approach to represent a user's information need in session search. [**Ch. 3**]
2. A Convolutional Neural Network (CNN) that learns to formulate queries that retrieve attachments in accordance to incoming request messages. [**Ch. 4**]
3. Three latent vector space models: the log-linear model for expert finding, Latent Semantic Entities (LSE) and the Neural Vector Space Model (NVSM). [**Ch. 5, 7 and 8**]

1.2.2 Theoretical contributions

4. A proactive email recommendation task, including a methodology for creating pseudo collections for model training and testing. [**Ch. 4**]
5. A formal framework for ranking attachments given a ranking over email messages that contain them. [**Ch. 4**]
6. Analyses of the time/space complexity of the log-linear model for expert finding and the NVSM. [**Ch. 5 and 8**]
7. A framework for analysing the quality of different entity representations, independent of the textual matching component. [**Ch. 7**]

1.2.3 Empirical contributions

8. (a) Analysis of the TREC Session Track that shows the prominence of short search sessions within the benchmarks. (b) Investigation of the effectiveness of specialized session search methods compared to our naive frequency-based approach. (c) Investigation of the viability of lexical query matching in session search. [Ch. 3]
9. Comparison of different query term formulation methods, their effectiveness and an analysis of the formulated queries and errors. [Ch. 4]
10. (a) Comparison of the log-linear model for expert finding with state-of-the-art retrieval methods, including traditional vector space models and language models. (b) Insight in how the uncertainty of the predictions of the log-linear model can be used to determine the effectiveness of the model. (c) Comparative error analysis between the semantic log-linear model and traditional generative language models that perform exact matching. (d) Insight in the relative strengths of semantic matching and exact matching for the expert retrieval task through an ensemble of the log-linear model and lexical language models. [Ch. 5]
11. Insight in the domain regularities (i.e., clusterings, similarity, importance) contained within latent entity representations (with an application to expert finding). [Ch. 6]
12. (a) A parameter sensitivity analysis of LSE models, with a focus on representation dimensionality and the amount of word context used to train the model. (b) Comparison of LSE models with state-of-the-art latent vector spaces in terms of retrieval effectiveness and according to the quality of the entity representations. (c) Insight in how LSE can benefit the retrieval performance in entity-oriented search engines that combine query-independent, lexical and semantic signals in a learning to rank model. [Ch. 7]
13. (a) Comparison of Neural Vector Space Model (NVSM) with lexical language models and state-of-the-art latent vector space models on article retrieval collections. (b) Analysis of the internals of NVSM and how it encodes word importance in the word representations. (c) Insight in the judgement bias inherent in TREC test collections. (d) Advice on how to configure the hyperparameters of NVSM. [Ch. 8]

1.2.4 Open-source software contributions

14. sesh (<https://github.com/cvangysel/sesh>) — a testbed for evaluating session search. [Ch. 3]
15. SERT (<https://github.com/cvangysel/SERT>) — the Semantic Entity Retrieval Toolkit that contains implementations of the log-linear model for expert finding and Latent Semantic Entities (LSE). [Ch. 5, 6, 7 and App. B]

16. cuNVSM (<https://github.com/cvangysel/cuNVSM>) — a highly-optimized CUDA implementation of LSE and NVSM that results in fast training and efficient memory usage. [Ch. 8]
17. pyndri (<https://github.com/cvangysel/pyndri>) — a Python interface to the Indri search engine. [Ch. 3, 4, 8 and App. A]

Only the software used for the experiments of Chapter 4 has not been released open-source as it is intellectual property of Microsoft Corporation.

1.3 Thesis overview

In this section we give an overview of the dissertation and provide recommendations for reading directions.

The chapter you are currently enjoying (Chapter 1) gives an introduction to the subject of this dissertation. In addition, the chapter provides an overview of the research questions and contributions. Chapter 2 discusses related work for the Chapters 3 to 8 that follow.

Part I of this dissertation contains research chapters related to query formulation from complex textual structures. In particular, Chapter 3 investigates the potential of lexical query formulation methods in session search. Chapter 4 introduces the task of proactive email attachment recommendation. In addition, it proposes a method that formulates a lexical query from an email thread

Part II of this dissertation introduces novel latent vector spaces for Information Retrieval. Chapter 5 targets a particular instance of the entity ranking task: expert finding. Chapter 6 performs additional analysis on the learned latent expert representations. The expert finding model is then adapted to larger entity domains (i.e., product search) in Chapter 7. Chapter 8 introduces NVSM, an extension to LSE that brings qualitative improvements. We evaluate NVSM on article retrieval and perform an in-depth analysis of its matching signal. In addition, we investigate the pool bias in off-line test collections and give practical advice on how to configure the hyperparameters of NVSM.

We conclude this dissertation in Chapter 9 and give directions for future work. Appendices A and B provide a description of some of the software that was developed as part of this dissertation.

Readers familiar with retrieval models and latent vector spaces may skip the appropriate parts of Chapter 2. Part I and Part II can be read independently of each other. If time is of the essence, only read Chapters 4, 8 and 9.

1.4 Origins

We list for each chapter the publications on which it is based. The dissertation is based on, in total, 9 publications [253, 256–263].

Chapter 3 is based on the conference paper *Lexical Query Modeling in Session Search* published at ICTIR’16 by Van Gysel, Kanoulas, and de Rijke [258].

The naive baseline method was designed by Van Gysel, experiments and analyses were performed by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Chapter 4 is based on the conference paper *Reply With: Proactive Recommendation of Email Attachments* published at CIKM'17 by Van Gysel, Mitra, Venanzi, Rosemarin, Kukla, Grudzien, and Cancedda [263].

The research was performed during a research internship at Microsoft Bing in London. The task was proposed by Mitra, Venanzi and Cancedda. The model was designed by Van Gysel and was inspired by ideas of Mitra. Some parts of the model were inspired through suggestions by Venanzi, Cancedda and Rosemarin. Kukla and Rosemarin helped by providing data. Grudzien performed additional analysis. Van Gysel did most of the writing, with the help of Mitra. Venanzi, Cancedda and Rosemarin also contributed to the text.

Chapter 5 is based on the conference paper *Unsupervised, Efficient and Semantic Expertise Retrieval* published at WWW'16 by Van Gysel, de Rijke, and Worring [253, 257].

The model was designed by Van Gysel, experiments and analyses were performed by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Chapter 6 is based on the conference paper *Structural Regularities in Expert Vector Spaces* published at ICTIR'17 by Van Gysel, de Rijke, and Kanoulas [259].

Experiments and analyses were performed by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Chapter 7 is based on the conference paper *Learning Latent Vector Spaces for Product Search* published at CIKM'16 by Van Gysel, de Rijke, and Kanoulas [256].

The model was designed by Van Gysel, experiments and analyses were performed by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Chapter 8 is based on the journal paper *Neural Vector Spaces for Unsupervised Information Retrieval* under review at TOIS by Van Gysel, de Rijke, and Kanoulas [260].

The model was designed by Van Gysel, experiments and analyses were performed by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Appendix A is based on the conference paper *Pyndri: a Python Interface to the Indri Search Engine* published at ECIR'17 by Van Gysel, Kanoulas, and de Rijke [262].

The software was implemented by Van Gysel. The software upon which Pyndri builds, Indri [238], was contributed by the Lemur project. All authors contributed to the text, Van Gysel did most of the writing.

Appendix B is based on the workshop paper *Semantic Entity Retrieval Toolkit* published at NeuIR’17 by Van Gysel, de Rijke, and Kanoulas [261].

The software was implemented by Van Gysel. All authors contributed to the text, Van Gysel did most of the writing.

Work performed as part of this dissertation also contributed to and benefited from insights gained through research that led to the following publications:

- C. Van Gysel. Listening to the flock - towards opinion mining through data-parallel, semi-supervised learning on social graphs. Master’s thesis, University of Antwerp, 2014
- C. Van Gysel, B. Goethals, and M. de Rijke. Determining the presence of political parties in social circles. In *ICWSM*, volume 2015, pages 690–693, 2015
- C. Van Gysel, L. Velikovich, I. McGraw, and F. Beaufays. Garbage modeling for on-device speech recognition. In *Interspeech*, volume 2015, pages 2127–2131, 2015
- C. Van Gysel, I. Oparin, X. Niu, and Y. Su. Rank-reduced token representation for automatic speech recognition, 2017. US Patent Application 15/459,481
- T. Kenter, A. Borisov, C. Van Gysel, M. Dehghani, M. de Rijke, and B. Mitra. Neural networks for information retrieval. In *SIGIR 2017*, pages 1403–1406. ACM, 2017

2

Background

In this chapter, we discuss the background for the research presented in this dissertation. The related work concerning Part I and Part II is covered in Section 2.1 and Section 2.2, respectively. We assume that the reader is familiar with the basic principles underlying modern information retrieval, as can be found in, e.g., [66, 165].

2.1 Query formulation

We first cover work related to application domains: session search (Section 2.1.1), proactive information retrieval (Section 2.1.2) and email (Section 2.1.3).

2.1.1 Session search

Many complex information seeking tasks, such as planning a trip or buying a car, cannot sufficiently be expressed in a single query [111]. These multi-faceted tasks are exploratory, comprehensive, survey-like or comparative in nature [207] and require multiple search iterations to be adequately answered [136]. Donato et al. [75] note that 10% of the user sessions of a web search engine (more than 25% of query volume) consists of such complex information needs.

The TREC Session Track [246] created an environment for researchers “to test whether systems can improve their performance for a given query by using previous queries and user interactions with the retrieval system.” The track’s existence led to an increasing number of methods aimed at improving session search. Yang et al. [281] introduce the Query Change Model (QCM), which uses lexical editing changes between consecutive queries in addition to query terms occurring in previously retrieved documents, to improve session search. They heuristically construct a lexicon-based query model for every query in a session. Query models are then linearly combined for every document, based on query recency [281] or document satisfaction [51, 155], into a session-wide lexical query model. However, there has been a clear trend towards the use of supervised learning [51, 157, 281] and external data sources [99, 156]. Guan et al. [99] perform lexical query expansion by adding higher-order n-grams to queries by mining document snippets. In addition, they expand query representations by including anchor texts to previously top-ranked documents in the session. Carterette et al. [51] expand document representations by including incoming anchor texts. Luo et al. [157]

introduce a linear point-wise learning-to-rank model that predicts relevance given a document and query change features. They incorporate document-independent session features in their ranker. The use of machine-learned ranking and the expansion of query and document representations is meant to address a specific instance of a wider problem in information retrieval, namely the query document mismatch [146]. In Chapter 3 of this dissertation, we analyse the session query logs made available by TREC and compare the performance of different lexical query modelling approaches for session search.

2.1.2 Proactive information retrieval

Zero-query search—or proactive IR—scenarios have received increasing attention recently [7]. However, similar approaches have also been studied in the past under other names, such as *just-in-time* [210–213], query-free [110] or anticipatory [46, 149] IR. According to Hart and Graham [110], the goal of the proactive retrieval system is to surface information that helps the user in a broader task. While some of these works focus on displaying contextually relevant information next to Web pages [46, 62, 163, 210, 212] or multimedia [195], others use audio cues [189, 225] or signals from other sensors [211, 220] to trigger the retrieval. In more recent years, proactive IR systems have re-emerged in the form of intelligent assistant applications on mobile devices, such as Siri, Google Now and Cortana. The retrieval in these systems may involve modelling repetitive usage patterns to proactively show concise information cards [233, 235] or surface them in response to change in user context such as location [31]. Hart and Graham [110], Budzik and Hammond [46] and Liebling et al. [149] propose to proactively formulate a query based on the user’s predicted information need.

In contrast to previous work on proactive contextual recommendation, we formulate a query to retrieve attachable items to assist users with composing emails instead of supplying information to support content or triggering information cards in mobile assistants. We propose a novel proactive retrieval task for email attachment recommendation in Chapter 4 of this dissertation.

2.1.3 Predictive models for email

Email overload is the inability to effectively manage communication due to the large quantity of incoming messages [275]. Grevet et al. [97] find that work email tends to be overloaded due to outstanding tasks or reference emails saved for future use. Ai et al. [5] find that 85% of email searches are targeted at retrieving known items (e.g., reference emails, attachments) in mailboxes. Horvitz [117] argues that a combination of two approaches—(1) providing the user with powerful tools, and (2) predicting the user’s next activity and taking automated actions on her behalf—is effective in many scenarios. Modern email clients may better alleviate email overload and improve user experience by incorporating predictive models that try to anticipate the user’s need and act on their behalf.

Missing attachments in email generates a wave of responses notifying the sender of her error. Dredze et al. [76] present a method that notifies the user when a file should be attached before the email is sent. Castro et al. [53] proposed a learning framework

to predict the action that will be performed on an email by the user, with the aim of prioritizing actionable emails. Carvalho and Cohen [52] classify emails according to their speech acts. Graus et al. [94], Qadir et al. [206] recommend recipients to send an email message to. Kannan et al. [127] propose an end-to-end method for automatically generating email responses that can be sent by the user with a single click. In Chapter 4 of this dissertation, we introduce a neural network architecture that learns to formulate a query to predict items to attach to an email reply.

2.1.4 Query formulation and reformulation

The task we study in Chapter 4, i.e., the task of contextual recommendation of attachable items by means of query formulation, has not received much attention. However, there is work on query extraction from verbose queries and query construction for related patent search. Similar to our work in Part I, the methods below consider the search engine as a black box.

2.1.4.1 Prior art search

Establishing novelty is an important part of the patenting process. Patent practitioners (e.g., lawyers, patent office examiners) employ a search strategy where they construct a query based on a new patent application in order to find prior art. However, patents are different from typical documents due to their length and lack of mid-frequency terms [158].

Automated query generation methods have been designed to help practitioners search for prior art. Xue and Croft [279] use TF-IDF to generate a ranking of candidate query terms, considering different patent fields, to rank similar patents. In later work [278], they incorporate a feature combination approach to further improve prior art retrieval performance. Alternative term ranking features, such as relative entropy [164], term frequency and log TF-IDF [54], have also been explored. Kim et al. [132] suggest boolean queries by extracting terms from a pseudo-relevant document set. Golestan Far et al. [93] find that an interactive relevance feedback approach outperforms state-of-the-art automated methods in prior art search.

2.1.4.2 Improving verbose queries

Bendersky and Croft [27] point out that search engines do not perform well with verbose queries [14]. Kumaran and Carvalho [140] propose a sub-query extraction method that obtains ground truth by considering every sub-query of a verbose query and cast it as a learning to rank problem. Xue et al. [280] use a Conditional Random Field (CRF) to predict whether a term should be included. However, inference using their method becomes intractable in the case of long queries. Lee et al. [143] learn to rank query terms instead of sub-queries with a focus on term dependency. Huston and Croft [119] find that removing the stop structure in collaborative question answering queries increases retrieval performance. Maxwell and Croft [167] propose a method that selects query terms based on a pseudo-relevance feedback document set. Meij et al. [173] identify semantic concepts within queries to suggest query alternatives. Related to the task of

improving verbose queries is the identification of important terms [287]. He and Ounis [113] note that the use of relevance scores for query performance prediction is expensive to compute and focus on a set of pre-retrieval features that are strong predictors of the query’s ability to retrieve relevant documents. Arguello et al. [10] apply query subset selection on spoken queries. Nogueira and Cho [192] apply reinforcement learning to query reformulation. See [104] for an overview on information retrieval with verbose queries.

In Chapter 3, we investigate how terms appearing in web search engines user sessions (e.g., query terms, terms in result page snippets) can be used to reformulate user queries in order to improve retrieval effectiveness. Chapter 4 introduces a neural network that formulates a query from email requests for the task of proactive attachment recommendation.

2.2 Latent semantic models

We first cover related work on Latent Semantic Models (LSM)—the subject of Part II—in Section 2.2.1, followed by the entity retrieval task—the application domain of Chapters 5 to 7—in Section 2.2.2. Finally, we review work from neural language modelling—which inspired the latent vector spaces introduced in this dissertation—in Section 2.2.3, and neural information retrieval.

2.2.1 Latent semantic models for information retrieval

The mismatch between queries and documents is a critical challenge in search [146]. Latent Semantic Models (LSMs) enable retrieval based on conceptual content, instead of exact term matches. Especially relevant to this dissertation is the class of unsupervised latent semantic models. We distinguish between count-based approaches (Section 2.2.1.1) and approaches where representations are learned using neural networks. The latter class of methods is split between methods that combine pre-trained word embeddings (Section 2.2.1.2) from neural language models (Section 2.2.3) and representations that are learned from scratch (Section 2.2.1.3) specifically for the task at hand.

2.2.1.1 Count-based approaches

Latent Semantic Indexing (LSI) [70] and probabilistic LSI (pLSI) [116] were introduced in order to mitigate the mismatch between documents and queries [146]. Blei et al. [38] proposed Latent Dirichlet Allocation (LDA), a topic model that generalizes to unseen documents.

2.2.1.2 Combining pre-trained embeddings

Vulić and Moens [270] are the first to aggregate word embeddings learned with a context-predicting distributional semantic model (DSM); query and document are represented as a sum of word embeddings learned from a pseudo-bilingual document collection with a Skip-gram model. Kenter and de Rijke [128] extract features from embeddings for the

task of determining short text similarity. Zuccon et al. [290] use embeddings to estimate probabilities in a translation model that is combined with traditional retrieval models (similar to [89, 247]). Zamani and Croft [282, 283] investigate the use of pre-trained word embeddings for query expansion and as a relevance model to improve retrieval. Guo et al. [103] introduce the Bag-of-Word-Embeddings (BoWE) representation where every document is represented as a matrix of the embeddings occurring in the document; their non-linear word transportation model compares all combinations of query/document term representations at retrieval time. They incorporate lexical matching into their model by exactly comparing embedding vector components for specific terms (i.e., specific terms occurring in both document and query are matched based on the equality of their vector components, contrary to their lexical identity).

2.2.1.3 Learning from scratch

The methods discussed above incorporate features from neural language models. The recent deep learning revival, however, was due to the end-to-end optimization of objectives and representation learning [137, 142, 239] in contrast to feature engineering or the stacking of independently-estimated models. The following neural methods learn representations of words and documents from scratch. Salakhutdinov and Hinton [223] introduce semantic hashing for the document similarity task. Le and Mikolov [141] propose doc2vec, a method that learns representations of words and documents. Ai et al. [4] evaluate the effectiveness of doc2vec representations for ad-hoc retrieval, but obtain dissappointing results that are further analysed in [3]. In Part II of this dissertation, we introduce three novel models that learn document representations from scratch.

2.2.2 Entity retrieval

Around 40% of web queries [202] concern entities. Entity-oriented queries express an information need that is better answered by returning specific entities as opposed to documents [22]. The entity retrieval task is characterized by a combination of (noisy) textual data and semi-structured knowledge graphs that encode relations between entities [74]. Entity and document retrieval [23, p. 224] are closely related as performance of the latter can greatly impact that of the former [160].

2.2.2.1 Expert retrieval

Early expert retrieval systems were often referred to as expert locator and expertise management systems [168]. These database systems typically relied on people to self-assess their expertise against a predefined set of topics [171], which is known to generate unreliable results [25].

With the introduction of the P@NOPTIC system [63], and later the TREC Enterprise track [245], there has been an active research interest in automated expertise profiling methods. It is useful to distinguish between *profile-based* methods, which create a textual representation of a candidate's knowledge, and *document-based* methods, which represent candidates as a weighted combination of documents. The latter generally perform better at ranking, while the former is more efficient as it avoids retrieving all documents relevant to a query [23, p. 221].

There has been much research on generative probabilistic models for expert retrieval [82, 200]. Such models have been categorized in candidate generation models [49], topic generation models [20, 21] and proximity-based variants [21, 228]. Of special relevance to us are the unsupervised profile-centric (Model 1) and document-centric (Model 2) models of Balog et al. [20], which focus on raw textual evidence without incorporating collection-specific information (e.g., query modelling, document importance or document structure). Supervised discriminative models [84, 187, 236] are preferred when query-candidate relevance pairs are available for training. Unlike their generative counterparts these models have no issue combining complex and heterogeneous features (e.g., link-based features, document importance features, etc.); they resemble Learning to Rank (L2R) methods for document retrieval [23, 151]. However, a lack of training data may greatly hinder their applicability [23, p. 179]. Beyond unsupervised generative and supervised discriminative approaches, there are graph-based approaches based on random walks [229] and voting-based approaches based on data fusion [159]. Demartini et al. [71] propose a vector space-based method for the entity ranking task; their framework extends vector spaces operating on documents to entities. Closely related to expert finding is the task of expert profiling, of which the goal is to describe an expert by her areas of expertise [18], and similar expert finding [18]; see [23] for an overview. In Chapter 5 of this dissertation we introduce a log-linear model that learns the relations between experts and words in an unsupervised manner from scratch. Compared to generative language models that perform term-based matching, our model contributes a complementary semantic matching signal. We show that an unsupervised example of the term-based lexical methods and our methods performs best.

2.2.2.2 Product retrieval

Product search engines are an important source of traffic in the e-commerce market [124]. Specialized solutions are needed to maximize the utilization of these platforms. Nurmi et al. [193] note a discrepancy between buyers' shopping lists and how retail stores maintain information. They introduce a grocery retrieval system that retrieves products using shopping lists written in natural language. Product resolution [16] is an important task for e-commerce aggregation platforms, such as verticals of major web search engines and price comparison websites. Duan et al. [79] propose a probabilistic mixture model for the attribute-level analysis of product search logs. They focus on structured aspects of product entities, while in this work we learn representations from unstructured documents. Duan et al. [80] extend the language modelling approach to product databases by incorporating the ability to condition on specification (e.g., lightweight products only). They note that while languages such as SQL can be used effectively to query these databases, their use is difficult for non-experienced end users. Duan and Zhai [78] study the problem of learning query intent representation for structured product entities. In Chapter 7 of this dissertation we introduce a latent vector space model that learns representations of entities, words and the connection in between. We show how our latent vector space contributes a complementary matching signal that can be incorporated in product search engines.

2.2.2.3 Representation learning

Part II of this dissertation covers the learning of entity representations. There already exists some work in this area. Bordes et al. [39] leverage structured relations captured in Knowledge Bases (KB) for entity representation learning and evaluate their representations on the link prediction task. Our approach has a strong focus on modelling the language of all entities collaboratively, without the need for explicit entity relations during training. Zhao et al. [289] employ matrix factorization methods to construct low-dimensional continuous representations of entities, categories and words for determining similarity of Wikipedia entities. They employ a word pair similarity evaluation set and only evaluate on pairs referring to Wikipedia entities; they learn a single semantic space for widely-differing concepts (entities, categories and words) of different cardinalities and make extensive use of an underlying Knowledge Graph (KG) to initialize their parameters.

In Part II of this dissertation, we learn low-dimensional representations of words, entities and documents. Entity representations are the subject of Chapter 5 (people) and Chapter 7 (products). The representations learned for people in Chapter 5 are analysed in Chapter 6. In Chapter 8, we learn representations of news article documents.

2.2.3 Neural language modelling

Large-vocabulary neural probabilistic language models for modelling word sequence distributions have become very popular recently [32, 181, 182]. These models learn continuous-valued distributed representations for words, also known as embeddings [177, 183, 199], in order to fight the curse of dimensionality and increase generalization by introducing the expectation that similar word vectors signify semantically or syntactically similar words. Neural Network Language Models (NNLM) [32, 175] have shown promising results in Natural Language Processing (NLP) [126, 239, 243] and Automatic Speech Recognition (ASR) [96, 222] compared to Markovian models. Collobert et al. [59] apply NNLMs to arbitrary NLP tasks by learning one set of word representations in a multi-task setting. Even more recently, there has been a surge in multi-modal neural language models [134], which lend themselves to the task of automated image captioning.

2.2.3.1 Representations and regularities

The idea that representations may capture linguistic or semantic regularities has received considerable attention. More generally, the idea of learning a representation of the elements of a discrete set of objects (e.g., words) is not new [33, 115, 219]. However, it has only been since the turn of the last century that NNLMs, which learn word embeddings as a side-effect of dealing with high-dimensionality, were shown to be better at modelling language than Markovian models [32, 181, 182].

Turian et al. [248] compare word representations learned by neural networks, distributional semantics and cluster-based methods as features in Named Entity Recognition (NER) and chunking. They find that both cluster-based methods and distributed word representations learned by NNLMs improve performance, although cluster-based methods yield better representations for infrequent words. Baroni et al. [24] confirm the

superiority of context-predicting (word embeddings) over context-counting (distributional semantics) representations.

Later algorithms are specifically designed for learning word embeddings [177, 183, 199], such that, somewhat ironically, NNLMs became a side-product. These embeddings contain linguistic regularities [144, 178], as evidenced in syntactic analogy and semantic similarity tasks. Multiple *word* representations can be combined to form *phrase* representations [176]. Clusterings of word embeddings can be used to discover word classes [176]. And insights gathered from word embedding algorithms can be used to improve distributional semantics [145].

2.2.4 Neural information retrieval

The recent revival of neural networks due to advances in computer vision [137], NLP [59, 177] and ASR [96] has led to an increasing interest in these technologies from the information retrieval community.

Beyond representation learning that we extensively discussed in Section 2.2.1, there are more applications of neural models in IR [65, 196]. In machine-learned ranking [151], we have RankNet [47]. In the class of supervised learning-to-match approaches, where clicks are available, there are DSSM [118, 232] and DSN [72]. Guo et al. [102] learn a relevance model by extracting features from BoWE representations in addition to corpus statistics such as inverse document frequency. Recently, Mitra et al. [179] have introduced a supervised document ranking model that matches using both local and distributed representations. Next to retrieval models there has been work on modelling user interactions with neural methods. Borisov et al. [40] introduce a neural click model that represents user interactions as a vector representation; in [41], they extend their work by taking into account click dwell time.

Part I

Query Formulation

3

Lexical Query Modelling in Session Search

3.1 Introduction

In Section 2.1.1, we discussed the expansion of query and document representations in the session search domain in order to bridge the vocabulary gap [146]. In this chapter, we analyse the session query logs made available by TREC and compare the performance of different lexical query modelling approaches for session search, taking into account session length.¹ In addition, we investigate the viability of lexical query models in a session search setting.

The main purpose of this chapter is to investigate the potential of lexical methods in session search and provide foundations for future research. We ask the following questions towards answering **RQ1**:

RQ1.1 Increasingly complex methods for session search are being developed, but how do naive methods perform?

RQ1.2 How well can lexical methods perform?

RQ1.3 Can we solve the session search task using lexical matching only?

3.2 Lexical matching for sessions

We define a search session s as a sequence of n interactions (q_i, r_i) between user and search engine, where q_i denotes a user-issued query consisting of $|q_i|$ terms $t_{i,1}, \dots, t_{i,|q_i|}$ and r_i denotes a result page consisting of $|r_i|$ documents $r_{i,1}, \dots, r_{i,|r_i|}$ returned by the search engine (also referred to as SERP). The goal, then, is to return a SERP r_{n+1} given a query q_{n+1} and the session history that maximizes the user's utility function.

In this chapter, we formalize session search by modelling an observed session s as a query model parametrized by $\theta^s = \{\theta_1^s, \dots, \theta_{|V|}^s\}$, where θ_i^s denotes the weight associated with term $t_i \in V$ (specified below). Documents d_j are then ranked in

¹An open-source implementation of our testbed for evaluating session search is available at <https://github.com/cvangysel/sesh>.

decreasing order of

$$\log P(d_j | s) = \sum_{k=1}^{|V|} \theta_k^s \log \theta_k^{d_j},$$

where θ^{d_j} is a lexical model of document d_j , which can be a language model (LM), a vector space model or a specialized model using hand-engineered features. Query model θ^s is a function of the query models of the interactions i in the session, θ^{s_i} (e.g., for a uniform aggregation scheme, $\theta^s = \sum_i \theta^{s_i}$). Existing session search methods [99, 281] can be expressed in this formalism as follows:

Term frequency (TF) Terms in a query are weighted according to their frequency in the query (i.e., $\theta_k^{s_i}$ becomes the frequency of term t_k in q_i). Queries q_i that are part of the same session s are then aggregated uniformly for a subset of queries. In this chapter, we consider the following subsets: the *first query*, the *last query* and the *concatenation of all queries* in a session. Using the last query corresponds to the official baseline of the TREC Session track [51].

Nugget [99] is a method for effective structured query formulation for session search. Queries q_i , part of session s , are expanded using higher order n-grams occurring in both q_i and snippets of the top- k documents in the previous interaction, $r_{i-1,1}, \dots, r_{i-1,k}$. This effectively expands the vocabulary by additionally considering n-grams next to unigram terms. The query models of individual queries in the session are then aggregated using one of the aggregation schemes. Nugget is primarily targeted at resolving the query-document mismatch by incorporating structure and external data and does not model query transitions. The method can be extended to include external evidence by expanding θ^s to include anchor texts pointing to (clicked) documents in previous SERPs.

Query Change Model (QCM) QCM [281] uses syntactic editing changes between consecutive queries in addition to query changes and previous SERPs to enhance session search. In QCM [281, Section 6.3], document model θ^d is provided by a language model with Dirichlet smoothing and the query model at interaction i , θ^{s_i} , in session s is given by

$$\theta_k^{s_i} = \begin{cases} 1 + \alpha(1 - P(t_k | r_{i-1,1})), & t_k \in q_{\text{theme}} \\ 1 - \beta P(t_k | r_{i-1,1}), & t_k \in +\Delta q \wedge t_k \in r_{i-1,1} \\ 1 + \epsilon \text{idf}(t_k), & t_k \in +\Delta q \wedge t_k \notin r_{i-1,1} \\ -\delta P(t_k | r_{i-1,1}), & t_k \in -\Delta q, \end{cases}$$

where q_{theme} are the session's theme terms, $+\Delta q$ ($-\Delta q$, resp.) are the added (removed) terms, $P(t_k | r_{i-1,1})$ denotes the probability of t_k occurring in SAT clicks, $\text{idf}(t_k)$ is the inverse document frequency of term t_k and $\alpha, \beta, \epsilon, \delta$ are parameters. The θ^{s_i} are then aggregated into θ^s using one of the aggregation schemes, such as the uniform aggregation scheme (i.e., the sum of the θ^{s_i}).

In Section 3.4, we analyse the methods listed above in terms of their ability to handle sessions of different lengths and contextual history.

Table 3.1: Overview of 2011, 2012, 2013 and 2014 TREC session tracks. For the 2014 track, we report the total number of sessions in addition to those sessions with judgements. We report the mean and standard deviation where appropriate; M denotes the median.

	2011	2012	2013	2014
Sessions				
Sessions	76	98	87	100 (1,021 total)
Queries per session	3.68 ± 1.79 ; M=3.00	3.03 ± 1.57 ; M=2.00	5.08 ± 3.60 ; M=4.00	4.34 ± 2.22 ; M=4.00
Unique terms per session	7.01 ± 3.28 ; M=6.50	5.76 ± 2.95 ; M=5.00	8.86 ± 4.38 ; M=8.00	7.79 ± 4.08 ; M=7.00
Topics				
Session per topic	1.23 ± 0.46 ; M= 1.00	2.04 ± 0.98 ; M= 2.00	2.18 ± 0.93 ; M= 2.00	20.95 ± 4.81 ; M= 21.00
Document judgements per topic	313.11 ± 114.63 ; M=292.00	372.10 ± 162.63 ; M=336.50	268.00 ± 116.86 ; M=247.00	332.33 ± 149.03 ; M=322.00
Collection				
Documents	21,258,800			15,702,181
Document length	$1,096.18 \pm 1,502.45$			$649.07 \pm 1,635.29$
Terms	3.40×10^7 (2.33×10^{10} total)		2.36×10^7 (1.02×10^{10} total)	
Spam scores	GroupX		Fusion	

3.3 Experiments

3.3.1 Benchmarks

We evaluate the lexical query modelling methods listed in Section 3.2 on the session search task (G1) of the TREC Session track from 2011 to 2014 [246]. We report performance on each track edition independently and on the track aggregate. Given a query, the task is to improve retrieval performance by using previous queries and user interactions with the retrieval system. To accomplish this, we first retrieve the 2,000 most relevant documents for the given query and then re-rank these documents using the methods described in Section 3.2. We use the “Category B” subsets of ClueWeb09 (2011/2012) and ClueWeb12 (2013/2014) as document collections. Both collections consist of approximately 50 million documents. Spam documents are removed before indexing by filtering out documents with scores (GroupX and Fusion, respectively) below 70 [60]. Table 3.1 shows an overview of the benchmarks and document collections.

3.3.2 Evaluation measures

To measure retrieval effectiveness, we report Normalized Discounted Cumulative Gain at rank 10 (NDCG@10) in addition to Mean Reciprocal Rank (MRR). The relevance judgements of the tracks were converted from topic-centric to session-centric according to the mappings provided by the track organizers.² Evaluation measures are then computed using TREC’s official evaluation tool, `trec_eval`.³

3.3.3 Systems under comparison

We compare the lexical query model methods outlined in Section 3.2. All methods compute weights for lexical entities (e.g., unigram terms) on a per-session basis, construct a structured Indri query [174] and query the document collection using `pyndri`.⁴ For fair comparison, we use Indri’s default smoothing configuration (i.e., Dirichlet smoothing with $\mu = 2500$) and uniform query aggregation for all methods (different from the smoothing used for QCM in [281]). This allows us to separate query aggregation techniques from query modelling approaches in the case of session search.

For Nugget, we use the default parameter configuration ($k_{\text{snippet}} = 10$, $\theta = 0.97$, $k_{\text{anchor}} = 5$ and $\beta = 0.1$), using the strict expansion method. We report the performance of Nugget without the use of external resources (RL2), with anchor texts (RL3) and with click data (RL4). For QCM, we use the parameter configuration as described in [157, 281]: $\alpha = 2.2$, $\beta = 1.8$, $\epsilon = 0.07$ and $\delta = 0.4$.

In addition to the methods above, we report the performance of an oracle that always ranks in decreasing order of ground-truth relevance. This oracle will give us an upper-bound on the achievable ranking performance.

²We take into account the mapping between judgements and actual relevance grades for the 2012 edition.

³https://github.com/usnistgov/trec_eval

⁴<https://github.com/cvangysel/pyndri>




Figure 3.1: Box plot of NDCG@10 on all sessions of the TREC Session track (2011–2014). The box depicts the first, second (median) and third quartiles. The whiskers are located at 1.5 times the interquartile range on both sides of the box. The square and crosses depict the average and outliers respectively.

3.3.4 Ideal lexical term weighting

We investigate the maximally achievable performance by weighting query terms. Inspired by Bendersky et al. [30], we optimize NDCG@10 for every session using a grid search over the term weight space. We sweep the weight of every term between -1.0 and 1.0 (inclusive) with increments of 0.1 , resulting in a total of 21 weight assignments per term. Due to the exponential time complexity of the grid search, we limit our analysis to the 230 sessions with 7 unique query terms or less (see Table 3.1). This experiment will tell us the maximally achievable retrieval performance in session search by the re-weighting of lexical terms only.

3.4 Results & discussion

In this section, we report and discuss our experimental results. Of special interest to us are the methods that perform lexical matching based on a user’s queries in a single session: QCM, Nugget (RL2) and the three variants of TF. Table 3.2 shows the methods’ performance on the TREC Session track editions from 2011 to 2014. No single method consistently outperforms the other methods. Interestingly enough, the methods based on term frequency (TF) perform quite competitively compared to the specialized session search methods (Nugget and QCM). In addition, the TF variant using all queries in a session even outperforms Nugget (RL2) on the 2011 and 2014 editions and QCM on nearly all editions. Using the concatenation of all queries in a session, while being an obvious baseline, has not received much attention in recent

Table 3.2: Overview of experimental results on 2011–2014 TREC Session tracks of the TF, Nugget and QCM methods (see Section 3.2). The ground-truth oracle shows the ideal performance (Section 3.3.3).

	2011		2012		2013		2014	
	NDCG@10	MRR	NDCG@10	MRR	NDCG@10	MRR	NDCG@10	MRR
Ground-truth oracle	0.777	0.868	0.695	0.865	0.517	0.920	0.410	0.800
TF (first query)	0.371	0.568	0.302	0.523	0.121	0.379	0.120	0.336
TF (last query)	0.358	0.598	0.316	0.586	0.133	0.358	0.156	0.458
TF (all queries)	0.448	0.685	0.348	0.604	0.162	0.477	0.174	0.478
Nugget (RL2)	0.437	0.677	0.352	0.609	0.163	0.488	0.173	0.476
Nugget (RL3)	0.442	0.678	0.360	0.619	0.162	0.488	0.172	0.477
Nugget (RL4)	0.437	0.677	0.352	0.609	0.163	0.488	0.173	0.476
QCM	0.440	0.661	0.342	0.575	0.160	0.484	0.162	0.450

Table 3.3: NDCG@10 for TF weighting (Section 3.2), ideal term weighting (Section 3.3.4) and the ground-truth oracle (Section 3.3.3).

	2011	2012	2013	2014
TF (all queries)	0.391	0.333	0.179	0.183
Ideal term weighing	0.589	0.528	0.361	0.296
Ground-truth oracle	0.716	0.682	0.593	0.453

literature or by TREC [246]. In addition, note that the best-performing (unsupervised) TF method achieves better results than the supervised method of Luo et al. [157] on the 2012 and 2013 tracks. Fig. 3.1 depicts the boxplot of the NDCG@10 distribution over all track editions (2011–2014). The term frequency approach using all queries achieves the highest mean/median overall. Given this peculiar finding, where a generic retrieval model performs better than specialized session search models, we continue with an analysis of the TREC Session search logs.

In Fig. 3.2 we investigate the effect of varying session lengths in the session logs. The distribution of session lengths is shown in the top row of Fig. 3.2. For the 2011–2013 track editions, most sessions consisted of only two queries. The mode of the 2014 edition lies at 5 queries per session. If we examine the performance of the methods on a per-session length basis, we observe that the TF methods perform well for short sessions. This does not come as a surprise, as for these sessions there is only a limited history that specialized methods can use. However, the TF method using the concatenation of all queries still performs competitively for longer sessions. This can be explained by the fact that as queries are aggregated over time, a better representation of the user’s information need is created. This aggregated representation naturally emphasizes important *theme terms* of the session, which is a key component in the QCM [281].

How do these methods perform as the search session progresses? Fig. 3.3 shows the performance of sessions of length five after every user interaction, when using all queries in a session (Fig. 3.3a) and when using only the previous query (Fig. 3.3b). We can see that NDCG@10 increases as the session progresses for all methods. Beyond half of the session, the session search methods outperform retrieving according to the last query in the session. We see that, for longer sessions, specialized methods (Nugget, QCM) outperform generic term frequency models. This comes as no surprise. Bennett et al. [35] note that users tend to reformulate and adapt their information needs based on observed results and this is essentially the observation upon which QCM builds.

Fig. 3.1 and Table 3.2 reveal a large NDCG@10 gap between the compared methods and the ground-truth oracle. How can we bridge this gap? Table 3.3 shows a comparison between frequency-based term weighting, the ideal term weighting (Section 3.3.4) and the ground-truth oracle (Section 3.3.3) for all sessions consisting of 7 unique terms or less (Section 3.3.4). Two important observations. There is still plenty of room for improvement using lexical query modelling only. Relatively speaking, around half of the gap between weighting according to term frequency and the ground-truth can be bridged by predicting better term weights. However, the other half of the performance gap cannot be bridged using lexical matching only, but instead requires a notion of semantic matching [146].

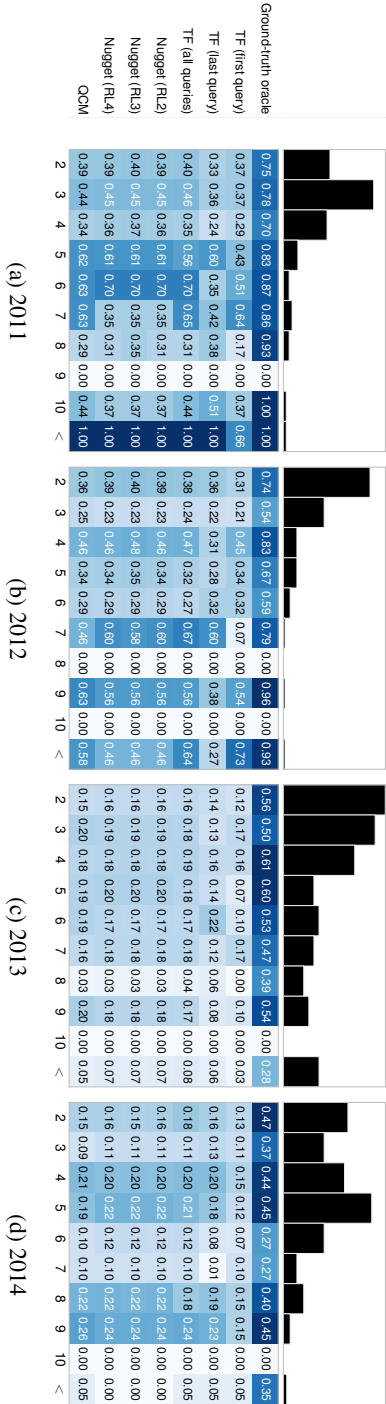


Figure 3.2: The top row depicts the distribution of session lengths for the 2011–2014 TREC Session tracks, while the bottom row shows the performance of the TF, Nugget and QCM models for different session lengths.




Figure 3.3: Difference in NDCG@10 with the official TREC baseline (TF using the last query only) of 5-query sessions (45 instances) with different history configurations for the 2011–2014 TREC Session tracks.

3.5 Summary

We have shown that naive frequency-based term weighting methods perform on par with specialized session search methods on the TREC Session track (2011–2014).⁵ This is due to the fact that shorter sessions are more prominent in the session query logs. On longer sessions, specialized models are able to exploit session history more effectively. Future work should focus on creating benchmarks consisting of longer sessions with complex information needs.

In the next chapter (Chapter 4), we introduce a neural network model that learns to formulate a query from complex textual structures (i.e., email threads). In Section 3.4, we observed that the query/document mismatch is prevalent in session search and methods restricted to lexical query modelling face a very strict performance ceiling. Therefore, Part II of this dissertation is dedicated to the modelling of latent vector spaces that bridge the vocabulary gap between query and document.

⁵An open-source implementation of our testbed for evaluating session search is available at <https://github.com/cvangysel/sesh>.

4

Reply With: Proactive Recommendation of Email Attachments

4.1 Introduction

In the previous chapter, we looked at formulating a query from user sessions in web search engines—a particular complex textual structure. In this chapter, we move our focus to formulating queries from a different complex textual structure: email threads. Email is still pervasive in the enterprise space [205], in spite of the growing popularity of social networks and other modern online communication tools. Users typically respond to incoming emails with textual responses. However, an analysis of the publicly available Avocado dataset [194] reveals that 14% of those messages also contain items, such as a file or a hyperlink to an external document. Popular email clients already detect when users forget to attach files by analysing the text of the response message [76, 77]. On Avocado, we find that in 35% of the cases where the response contains attachments, the item being attached is also present in the sender’s mailbox at the time of composing the response. This implies that modern email clients could help users compose their responses faster by proactively retrieving and recommending relevant items that the user may want to include with their message.

In *proactive* information retrieval (IR) systems [31, 149, 233, 235], the user does not initiate the search. Instead, retrieval is triggered automatically based on a user’s current context. The context may include the time of day [235], the user’s geographic location [31], recent online activities [233] or some other criteria. In our scenario, retrieval is based on the context of the current conversation, and in particular, the message the user is responding to. In a typical IR scenario, items are ranked based on query-dependent feature representations. In the absence of an explicit search query from the user, proactive IR models may formulate a keyword-based search query using the available context information and retrieve results for the query using a standard IR model [149]. Search functionalities are available from most commercial email providers and email search has been studied in the literature [5]. Therefore, we cast the email attachment recommendation problem as a query formulation task and use an existing IR system to retrieve emails. Attachable items are extracted from the retrieved emails and a ranking is presented to the user.

Fig. 4.1 shows an example of an email containing an explicit request for a file. In

4. Reply With: Proactive Recommendation of Email Attachments

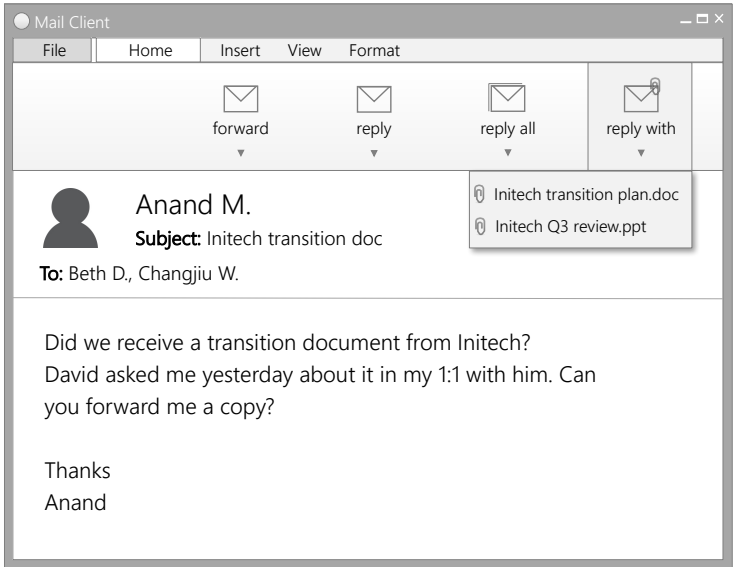


Figure 4.1: Anand asks Beth and Changjiu to forward him a copy of the Initech¹ transition document. Beth’s email client recommends two files, part of earlier emails in Beth’s mailbox, for her to attach to her reply.

general, there may or may not be an explicit request, but it may be appropriate to attach a relevant file with the response. Our task is to recommend the correct “transition document” as an attachment when Beth or Changjiu is responding to this email. In order to recommend an attachment, the model should formulate a query, such as “*Initech transition*”, based on the context of the request message, that retrieves the correct document from Beth’s or Changjiu’s mailbox. To formulate an effective query, the model must identify the discriminative terms in the message from Anand that are relevant to the actual file request.

Machine learning models that aim to solve the query formulation task need reliable feedback on what constitutes a good query. One option for generating labeled data for training and evaluation involves collecting manual assessments of proposed queries or individual query terms. However, it is difficult for human annotators to determine the ability of a proposed query to retrieve relevant items given only the request message. In fact, the efficacy of the query depends on the target message that should be retrieved, as well as how the IR system being employed functions. The relevance of the target message, in turn, is determined by whether they include the correct item that should be attached to the response message. Therefore, instead we propose an evaluation framework that requires an email corpus but no manual assessments. Request/response message pairs are extracted from the corpus and the model, that takes the request message as input, is evaluated based on its ability to retrieve the items attached to the

¹ Initech is a fictional company from a popular 1990s comedy film. Any resemblance to real organizations is purely coincidental.

response message. An IR system is employed for the message retrieval step, but is treated as a *black box* in the context of evaluation. Our framework provides a concise specification for the email attachment recommendation task (Section 4.3).

Our proposed approach for training a deep convolutional neural network (CNN) for the query formulation step is covered in Section 4.4. The model predicts a distribution over all the terms in the request message and terms with high predicted probability are selected to form a query. Model training involves generating a dataset of request/attachment pairs similar to the case of evaluation. Candidate queries are algorithmically synthesized for each request/attachment pair such that a message from the user’s mailbox with the correct item attached is ranked highly. We refer to synthetic queries as the *silver-standard queries* (or silver queries for brevity) to emphasize that they achieve reasonable performance on the task, but are potentially sub-optimal. The neural model is trained to minimize the prediction loss w.r.t. the silver queries given the request message as input.

The research questions we ask in this chapter towards answering **RQ1** are as follows:

RQ1.4 Do convolutional neural networks (CNN) improve ranking efficacy over state-of-the-art query formulation methods?

RQ1.5 When do CNNs work better than non-neural methods on the attachable item recommendation task?

RQ1.6 What features are most important when training CNNs?

4.2 Related work

We refer to Section 2.1 of our background chapter (Chapter 2) where the sections on proactive information retrieval (Section 2.1.2) and predictive models in email (Section 2.1.3) provide background on the problem domain of this chapter. Section 2.1.4 provides background on query formulation methods, with an emphasis on prior art search and sub-query selection in ad-hoc document retrieval. We elaborate briefly on the similarities and differences between the email domain and the domains of prior art and ad-hoc document retrieval.

Query extraction methods used for prior art search (Section 2.1.4.1) can also be applied to the task of attachable item recommendation considered in this chapter. Consequently, we consider the methods mentioned above as our baselines (Section 4.5.4). However, there are a few notable differences between the patent and email domains: (1) Email messages are much shorter in length than patents. (2) Patents are more structured (e.g., US patents contain more than 50 fields) than email messages. (3) Patents are linked together by a static citation graph that grows slowly, whereas email messages are linked by means of a dynamic conversation that is fast-paced and transient in nature. (4) In the case of email, there is a social graph between email users that can act as an additional source of information.

The task we consider in this chapter differs from search sub-query selection (Section 2.1.4.2) as follows. (1) Search queries are formulated by users as a way to interface with a search engine. Requests in emails may be more complex as they are formulated to retrieve information from a human recipient, rather than an automated search engine.

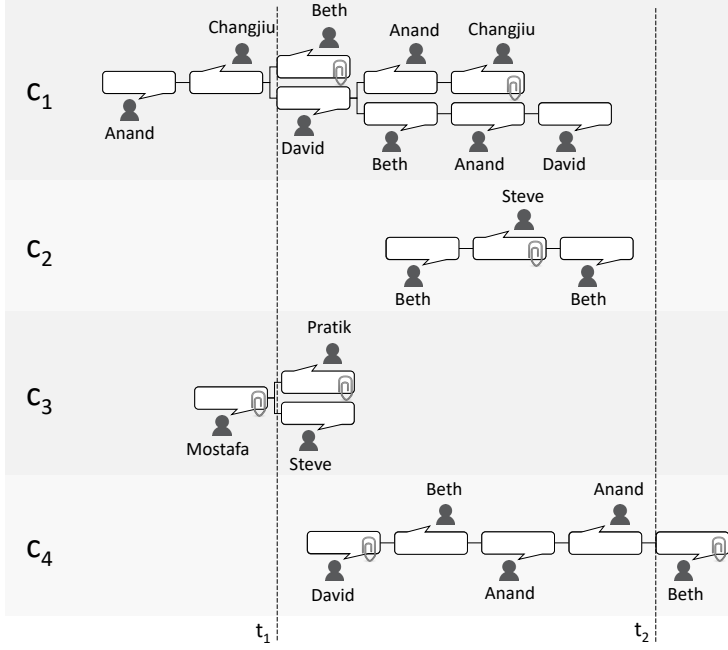


Figure 4.2: Beth’s mailbox has four on-going conversations. When Beth replies with an attachment in conversation c_1 , at time t_1 , only the attachment she received from Mostafa (conversation c_3) is present in her mailbox.

In other words, email requests are more likely to contain natural language and figurative speech than search engine queries. This is because the sender of the request does not expect their message to be parsed by an automated system. (2) Search sub-query extraction aims to improve retrieval effectiveness while the query intent remains fixed. This is not necessarily the case in our task, as a request message like has the intent to retrieve information from the recipient (rather than a retrieval system operating on top of the recipient’s mailbox). (3) Work on search sub-query selection [140, 280] takes advantage of the fact that 99.9% of search queries consist of 12 terms or less [28] by relying on computations that are intractable otherwise. As emails are longer (Table 4.2), many of the methods designed for search sub-query selection are not applicable in our setting.

4.3 Proactive attachable item recommendation

Given message $m_{a \rightarrow b}$ from user u_a to user u_b , we want to recommend an item e that the receiver u_b may want to attach (or include) in the response $m_{b \rightarrow a}$. Email corpora, such as Avocado [194], contain many conversation threads where each conversation c contains messages $m_i \in c$ exchanged between several participants. From these conversations, we can identify pairs of request-response messages $\langle m_{a \rightarrow b}, m_{b \rightarrow a} \rangle$

where $m_{b \rightarrow a}$ contains an attachment e_{actual} . We assume that user u_b included e_{actual} in $m_{b \rightarrow a}$ in response to an explicit or an implicit request in the message $m_{a \rightarrow b}$. Such pairs of request message and attachment $\langle m_{a \rightarrow b}, m_{b \rightarrow a} \rangle$ form the ground-truth in our evaluation framework.

Fig. 4.2 shows a sample mailbox M_{Beth} of user u_{Beth} containing four conversations $\{c_1, c_2, c_3, c_4\}$. During these conversations, u_{Beth} responds with an attachment twice—at time t_1 and t_2 . At time t_1 , in this toy example, the set of candidate items that are available in the user’s mailbox for recommendation contains only the attachment from $u_{Mostafa}$ received during the conversation c_3 . At t_2 , however, the set of candidates includes attachments received on all four conversation threads—from $u_{Changjiu}$ (c_1), u_{Steve} (c_2), $u_{Mostafa}$ (c_3), u_{Pratik} (c_3), and u_{David} (c_4)—as well as the item sent by u_{Beth} previously on the conversation thread c_1 .

It is important to emphasize that our problem setting has two important constraints when recommending items, that any model should adhere to (1) a **privacy** constraint: the model can only recommend items from a user’s own mailbox, and (2) a **temporal** constraint: the model can only recommend items that are already present in the user’s mailbox at the time of recommendation.

4.3.1 Attachment retrieval

In addition to the above domain-specific constraints, we limit our setup to using a standard IR system R for retrieval, and cast the problem that the model needs to solve as a query formulation task. Using an existing IR system has the practical benefit that one only needs to maintain a single system in contrast to the alternative where a separate attachment recommendation engine needs to be maintained. The model is presented with a message m_{req} containing an explicit or an implicit content request. The model is tasked with generating a query that can be submitted to the retrieval system R that retrieves a set of ranked messages M_R from the user’s mailbox. Under this assumption, the retrieval system R is treated as a black box, and we are only interested in optimizing the query formulation model. Note that a query is only formulated when it is clear that an item needs to be attached to a reply message (Section 4.5.1), such as is the topic of [76, 77].

To extract a ranked list of attachable items from search engine R , we adopt an approach popular in entity retrieval frameworks [20] where an entity model is the mixture of document models that the entity is associated with. For a given query q issued at time t' in the mailbox of user u , attachable items $e \in E$ are then ranked in decreasing order of

$$P(e | q, u, t') \propto \frac{1}{Z_1(e, u, t')} \sum_{\substack{m \in M \\ t_m < t'}} S_R(m | q) f(e | m) \quad (4.1)$$

where $S_R(m | q)$ is the relevance score for message m given query q according to retrieval system R , t_m is the timestamp when the message m appeared first in the mailbox M , t' is time when the model needs to make the recommendation, $f(e | m)$

denotes the association strength between message m and item e , and

$$Z_1(e, u, t') = \sum_{\substack{m \in M \\ t_m < t'}} f(e | m)$$

is a normalization constant. The normalization constant $Z_1(e, u, t')$ avoids a bias towards attachable items that are associated with many messages (e.g., electronic business cards).

We associate messages with an attachable item according to the presence of the item within a message and its surrounding messages within the conversation: $f(e | m) = \mathbb{1}_{\text{context}(m)}(e)$. In this chapter, we take $\text{context}(m)$ to be all messages m' in the same conversation c_m as message m that occurred before the time of recommendation, i.e., $t_{m'} < t'$. Note that the exact definition of an attachable item depends on the email domain and can include individual files, file bundles and hyperlinks to external documents amongst others (see Section 4.5.3).

4.3.2 Evaluating query formulations

Once we have extracted $\langle m_{\text{req}}, e_{\text{actual}} \rangle$ pairs from an email corpus, each request message m_{req} is presented to the query formulation model that we want to evaluate. The model generates a query q conditioned on the message m_{req} . The query q is submitted to retrieval system R and attachable items extracted from the retrieved messages are determined according to Eq. 4.1. Given the ranked list of attachable items $E_{\text{retrieved}}$ and the expected item e_{actual} we can compute standard rank-based IR metrics such as MRR and NDCG (Section 4.5.5). We report the mean metric over all $\langle m_{\text{req}}, e_{\text{actual}} \rangle$ pairs extracted from the corpus.

Our approach of using the historical information from an email corpus for evaluation is comparable to the application of *click-through* data for similar purposes in Web search. In the document retrieval scenario, a user's click on a document d on the search result page is considered an implicit vote of confidence on the relevance of d to the query. Learning to rank models can be trained on this *click-through* data [125, 161, 277] if human relevance judgements are not available in adequate quantity. By explicitly attaching a file e_{actual} , similarly, the user of an email system provides a strong indication that recommending e_{actual} at the time of composing m_{res} would have been useful. We can use this information to train and evaluate supervised models for ranking attachments at the time of email composition.

4.4 Query formulation model

We first introduce a method for generating pseudo training data [11, 12, 26, 37, 120, 121, 131, 240, 241] without the need for manual assessments. Silver-standard queries are algorithmically synthesized for each request/attachment pair and consequently scored by measuring the query's ability to retrieve the relevant attachment (Section 4.4.1). The request/query pairs part of the pseudo training collection are then used to train a convolutional neural network (Section 4.4.2) that learns to extract query terms from

a request message. In this chapter, we use a convolutional architecture rather than a recurrent one, as we intend to model term importance by term context without relying on the exact ordering of terms.

4.4.1 Model training

The candidate silver queries are extracted for request-response pairs $\langle m_{\text{req}}, m_{\text{res}} \rangle$ in a training set. Given a request message m_{req} and its associated target attachable item $e_{\text{actual}} \in E_{m_{\text{res}}}$ that is attached to reply m_{res} , where $E_{m_{\text{res}}}$ is the set of items attached to m_{res} , the objective is to select the k terms that are most likely to retrieve item e_{actual} according to Eq. 4.1.

In the ideal case, one considers the powerset of all terms within request message m_{req} as candidate silver queries [140, 280]. However, considering all terms is computationally intractable in our case as email messages tend to average between 70 to 110 tokens (Table 4.2).

In order to circumvent the intractability accompanied with computing the powerset of all terms in a message, we use the following stochastic strategy to select a fixed number of candidate query terms that we compute the powerset of. We consider two sources of query terms. The first source of candidate query terms consists of **subject** terms: topic terms in the subject of the request message. Email subjects convey relevance and context [274] and can be seen as a topical summary of the message. For the second source of query terms, we consider **recallable** terms: infrequent terms that occur frequently in messages $M_{\text{recallable}}(e_{\text{actual}}, t') = \{m \in M \mid t_m < t', e_{\text{actual}} \in E_m\}$ that contained item e_{actual} and occur at least once in the request message.. That is, we gather all terms that have the potential to retrieve e_{actual} (according to Eq. 4.1) and select those terms that occur in at least 30% of messages $M_{\text{recallable}}(e_{\text{actual}}, t')$ and occur in less than 1% of all messages.

To construct candidate silver queries for a request message m_{req} , we follow the strategy as outlined in Algorithm 1 that mimics the boolean query formulation process of Salton et al. [224]. Candidate terms are selected from either the **subject** or **recallable** source in increasing order of document frequency (i.e., infrequent terms first). Unwanted terms, such as stopwords, digits, punctuation and the names of the sender and recipients that occur in the email headers, are removed. Afterwards, we take the candidate queries $\tilde{Q}_{m_{\text{req}}}$ to be all possible subsets of candidate terms (excluding the empty set).

Once we have obtained the set of candidate queries $\tilde{Q}_{m_{\text{req}}}$ for request message m_{req} we score the candidate queries as follows. For every $\tilde{q} \in \tilde{Q}_{m_{\text{req}}}$ we rank email messages using retrieval system R according to \tilde{q} . We then apply Eq. 4.1 to obtain a ranking over items $E_{u, t'}$ in the mailbox of user u at time t' . As we know the target item e_{actual} to be retrieved for request message m_{req} , we quantify the performance of candidate query \tilde{q} by its reciprocal rank, $\text{score}(\tilde{q}) = \frac{1}{\text{rank}(e_{\text{actual}})} \in (0, 1]$ where $\text{rank}(e_{\text{actual}})$ denotes the position of item e_{actual} (Eq. 4.1) in the item ranking.

After computing the score for every candidate silver query, we group queries that perform at the same level (i.e., that have the same score) for a particular request message m_{req} . We then apply two post-processing steps that improve silver-standard query quality based on the trade-off between query broadness and specificity. Following

Algorithm 1: Candidate query terms are selected by choosing a random query term source and selecting the query term with the lowest document frequency while ignoring unwanted query terms [224]. The `isUnwanted` predicate is true when the term is a stopword, contains a digit, contains punctuation or equals the names of the email sender/recipients. After selecting k terms, we consider the powerset of selected terms as silver queries.

Data: request message m_{req} , query term budget k

Result: candidate silver queries $\tilde{Q}_{m_{\text{req}}}$

set of candidate terms $T \leftarrow \{\}$;

while (*term candidates left* $\wedge |T| < k$) **do**

$S \leftarrow$ uniformly random choose **subject** or **recallable terms**;

if $S = \emptyset$ **then**

continue

$t \leftarrow \arg \min_{t \in S} \text{df}(t)$;

if $\neg \text{isUnwanted}(t)$ **then**

$T = T \cup \{t\}$;

$S \leftarrow S \setminus \{t\}$

$\tilde{Q}_{m_{\text{req}}} \leftarrow 2^T - \{\emptyset\}$

Salton et al. [224] on boolean query formulation, specific queries are preferred over broad queries to avoid loss in precision. Queries can be made more specific by adding terms. Consequently, within every group of equally-performing queries, we remove subset queries whose union results in another query that performs at the same level as the subsets. For example, if the queries “*barack obama*”, “*obama family*” and “*barack obama family*” all achieve the same reciprocal rank, then we only consider the latter three-term query and discard the two shorter, broader queries. An additional argument for the strategy above follows from the observation that any term not part of the query is considered as undesirable during learning. Therefore, including all queries listed above as training material would introduce a negative bias against the terms “*barack*” and “*family*”. However, queries that are too specific can reduce the result set [224] or cause query drift [180]. Therefore, the second post-processing step constitutes the removal of supersets of queries that perform equal or worse. The intuition behind this is that the inclusion of the additional terms in the superset query did not improve retrieval performance. For example, if queries “*barack obama*” and “*barack obama president*” perform equally well, then the addition of the term “*president*” had no positive impact on retrieval. Consequently, including the superset query (i.e., “*barack obama president*”) in the training set is likely to motivate the inclusion of superfluous terms that negatively impact retrieval effectiveness.

Rank	Term	Score
1.	initech	0.20
2.	initech	0.18
3.	transition	0.15
4.	-----EoR-----	0.10
5.	david	0.07
6.	...	

Figure 4.3: Terms in the request message of Fig. 4.1 are ranked by our model. In addition to the terms, the model also ranks a (EoR) token that specifies the query end. The final query becomes “*initech transition*” as duplicate terms are ignored.

4.4.2 A convolutional neural network for ranking query terms

After obtaining a set of candidate queries $\tilde{Q}_{m_{\text{req}}}$ for every request/item pair $\langle m_{\text{req}}, e_{\text{actual}} \rangle$ in the training set, we learn to select query terms from email threads using a convolutional neural network model that convolves over the terms contained in the email thread. Every term is characterized by its context and term importance features that have been used to formulate queries in previous work [27, 54, 113, 140, 164, 278, 288]. Our model jointly learns to (1) generate a ranking of message terms, and (2) determine how many terms of the message term ranking should be included in the query. In order to determine the number of terms included in the query, the model learns to rank an end-of-ranking token (EoR) in addition to the message terms. Fig. 4.3 shows a term ranking for the example in Fig. 4.1. Terms in the request message are ranked in decreasing order of the score predicted by our model. Terms appearing at a lower rank than the (EoR) are not included in the query.

Our convolutional neural network (CNN) term ranking model is organized as follows; see Fig. 4.4 for an overview. Given request message m_{req} , we perform a convolution over the n message terms w_1, \dots, w_n . Every term w_k is characterized by (1) the term w_k itself, (2) the $2 \cdot L$ terms, $w_{k-L}, \dots, w_{k-1}, w_{k+1}, \dots, w_{k+L}$, surrounding term w_k where L is a context width hyperparameter, and (3) auxiliary query term quality features (see Table 4.1).

For every term in the message, the local context features (1st part of Table 4.1) are looked up in term embedding matrix W_{repr} (learned as part of the model) and the auxiliary features (part 2-4 of Table 4.1) are computed. For the auxiliary features, we apply min-max feature scaling on the message-level such that they fall between 0 and 1. The flattened embeddings, concatenated with the auxiliary feature vector, are fed to the neural network. At the output layer, the network predicts a term ranking score, $g(w_k, m_{\text{req}})$, for every term. In addition, a score for the (EoR) token, $h(m_{\text{req}})$, is predicted as well. The (EoR) score function h takes the same form as the term score function g , but has a separate set of parameters and takes as input an aggregated vector that represents the whole message. More specifically, the input to the (EoR) score function is the average of the term representations and their auxiliary features.

The ranking scores are then transformed into a distribution over message terms and

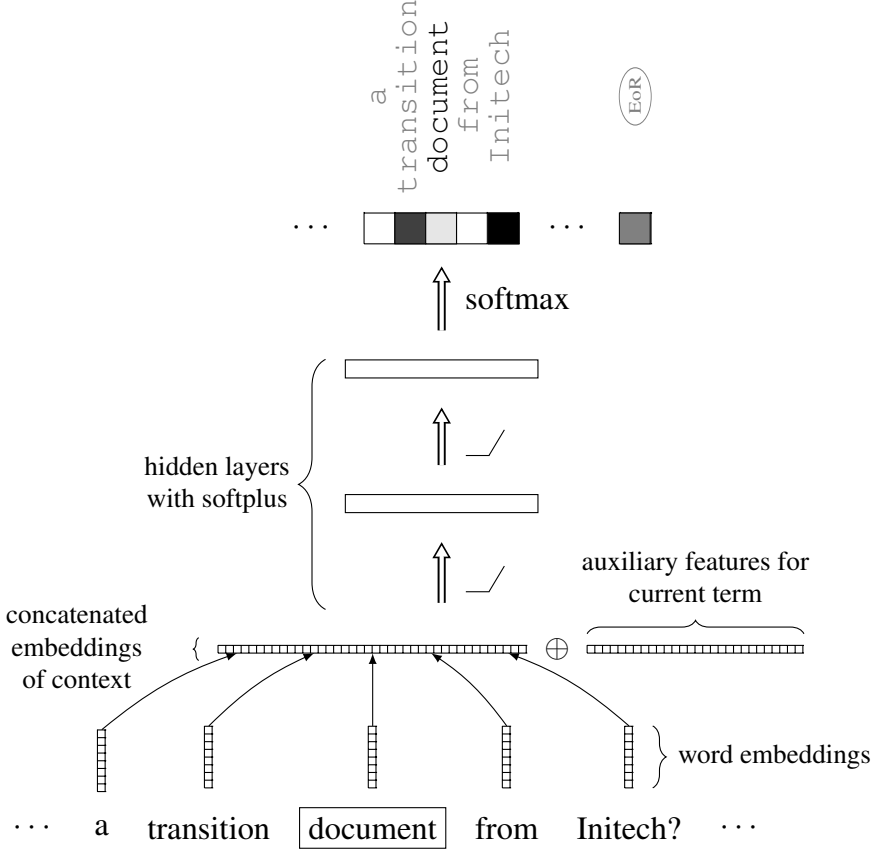


Figure 4.4: The model convolves over the terms in the message. For every term, we represent it using the word representations of its context. These representations are learned as part of the model. After creating a representation of the term’s context, we concatenate auxiliary features (Table 4.1). At the output layer, a score is returned as output for every term in the message. The softmax function converts the raw scores to a distribution over the message terms and the (EoR) token. Grayscale intensity in the distribution depicts probability mass.

the (EoR) token as follows:

$$P(w_k | m_{\text{req}}) = \frac{1}{Z_2} e^{g(w_k, m_{\text{req}})}$$

$$P(\text{(EoR)} | m_{\text{req}}) = \frac{1}{Z_2} e^{h(m_{\text{req}})}$$

with $Z_2(m_{\text{req}}) = e^{h(m_{\text{req}})} + \sum_{l=1}^{|m_{\text{req}}|} e^{g(w_l, m_{\text{req}})}$ as a normalization constant. For every

Table 4.1: Overview of term representation (learned as part of the model) and auxiliary features.

Context features (learned representations)	
term	Representation of the term.
context	Representations of the context surrounding the term.
Part-of-Speech features	
is_noun	POS tagged as a noun [27]
is_verb	POS tagged as a verb
is_other	POS tagged as neither a noun or a verb
Message features	
is_subject	Term occurrence is part of the subject [54]
is_body	Term occurrence is part of the body [54]
Abs. TF	Abs. term freq. within the message [278]
Rel. TF	Rel. term freq. within the message [278]
Rel. pos.	Rel. position of the term within the message
is_oov_repr	Term does not have a learned representation
Collection statistics features	
IDF	Inverse document frequency of the term [278]
TF-IDF	TF \times IDF [278]
Abs. CF	Abs. collection freq. within the collection
Rel. CF	Rel. collection freq. within the collection
Rel. Entropy	KL divergence from the unsmoothed collection term distribution to the smoothed ($\lambda = 0.5$) document term distribution [164]
SCQ	Similarity Collection/Query [288]
ICTF	Inverse Collection Term Frequency [140]
Pointwise SCS	Pointwise Simplified Clarity Score [113]

query $\tilde{q} \in \tilde{Q}$, the ground-truth distribution equals:

$$\begin{aligned}
 Q(w_k | \tilde{q}) &= \alpha \cdot \frac{\mathbb{1}_{\tilde{q}}(w_k)}{\#(w_k, m_{\text{req}}) \cdot |\tilde{q}|} \\
 Q(\text{EoR} | \tilde{q}) &= (1 - \alpha)
 \end{aligned} \tag{4.2}$$

where $\alpha = 0.95$ is a hyperparameter that determines the probability mass assigned to the EoR token and $\mathbb{1}_{\tilde{q}}(w_k)$ is the indicator function that evaluates to 1 when term w_k is part of silver query \tilde{q} . The frequency count $\#(w_k, m_{\text{req}})$ denotes the number of times term w_k occurs in message m_{req} and is included such that frequent and infrequent message terms are equally important.

Eq. 4.2 assigns an equal probability to every unique term in message m_{req} that occurs in silver query \tilde{q} . Our cost function consists of two objectives. The first objective aims to make $P(\cdot | m_{\text{req}})$ close to $Q(\cdot | \tilde{q})$ by minimizing the cross entropy:

$$L_{\text{xent}}(\theta | m_{\text{req}}, \tilde{q}) = - \sum_{\omega \in \Omega} Q(\omega | \tilde{q}) \log(P(\omega | m_{\text{req}})) \tag{4.3}$$

where $\Omega = (w_1, \dots, w_n, \text{EoR})$ is the sequence of all terms in the message m_{req}

concatenated with the end-of-ranking token. Eq. 4.3 promotes term ranking precision as it causes terms in the silver query to be ranked highly, immediately followed by the end-of-ranking token. The second objective encourages term ranking recall by dictating that the (EoR) token should occur at the same rank as the lowest-ranked silver query term:

$$L_{\text{cutoff}}(\theta \mid m_{\text{req}}, \tilde{q}) = (\min_{w \in \tilde{q}} (g(w, m_{\text{req}})) - h(m_{\text{req}}))^2 \quad (4.4)$$

The two objectives (Eq. 4.3-4.4) are then combined in a batch objective:

$$\begin{aligned} L(\theta \mid B) = & \frac{1}{|B|} \sum_{(m, \tilde{q}) \in B} \text{score}(\tilde{q})(L_{\text{xent}}(\theta \mid m, \tilde{q}) + L_{\text{cutoff}}(\theta \mid m, \tilde{q})) \\ & + \frac{1}{2\lambda} \sum_{W \in \theta_W} \sum_{ij} W_{ij}^2 \end{aligned} \quad (4.5)$$

where B is a uniformly random sampled batch of message/query pairs, θ_W is the set of parameter matrices and λ is a weight regularization parameter. Objective 4.3 resembles a list-wise learning to rank method [50] where a softmax over the top-ranked items is used. Eq. 4.5 is then optimized using gradient descent.

4.5 Experimental set-up

4.5.1 Research questions

As indicated in the introduction of this chapter, we seek to answer the following research questions:

RQ1.4 Do convolutional neural networks (CNN) improve ranking efficacy over state-of-the-art query formulation methods?

What if we consider the different fields (subject and body) in the email message when selecting query terms? To what extent do methods based on selecting the top ranked terms according to term scoring methods (e.g., TF-IDF, RE) perform? Can CNNs outperform state-of-the-art learning to rank methods? What can we say about the length of the queries extracted by the different methods?

RQ1.5 When do CNNs work better than non-neural methods on the attachable item recommendation task?

In the case that CNNs improve retrieval effectiveness over query extraction methods: what can we say about the errors made by CNNs? In particular, in what cases do our deep convolutional neural networks perform better or worse compared to the query term ranking methods under comparison?

RQ1.6 What features are most important when training CNNs?

Are all types of features useful? Can we make any inferences about the email domain or the attachable item recommendation task?

Table 4.2: Overview of the enterprise email collections used in this chapter: Avocado (public) and PIE (proprietary).

	Avocado	PIE
Messages	928,992	1,047,311
Message length (terms)	112.33 ± 244.01	74.70 ± 551.88
Threads	804,010	381,448
Thread lengths	1.19 ± 0.70	2.75 ± 3.65
Time period	3 years, 8 months	1 year
Attachable entities	50,462	28,725
Impressions per item	3.48 ± 2.55	2.79 ± 1.36
Messages with an item	311,478	152,649
no thread history	288,099	69,796
all items filtered (Section 4.5.3)	22,399	80,717
Request/reply pairs	980	2136
Thread history length of pairs	1.53 ± 1.13	4.04 ± 5.78
Relevant items per pair	1.22 ± 0.70	1.29 ± 1.82

4.5.2 Experimental design

We operate under the assumption that an incoming message has been identified as a request for content. A query is then formulated from the message using one of the query formulation methods (Section 4.5.4). To answer the research questions posed in Section 4.5.1, we compare CNNs with existing state-of-the-art query term selection methods on enterprise email collections (**RQ1.4**). In addition, we look at the query lengths generated by the formulation methods that perform best. **RQ1.5** is answered by examining the per-instance difference in Reciprocal Rank (RR) (Section 4.5.5). After that, we perform a qualitative analysis where we examine the outlier examples. For **RQ1.6** we perform a feature ablation study where we systematically leave out a feature category (Table 4.1).

4.5.3 Data collections and pre-processing

We answer our research questions (Section 4.5.1) using two enterprise email collections that each constitute a single tenant (i.e., an organization): (1) the Avocado collection [194] is a public data set that consists of emails taken from 279 custodians of a defunct information technology company, and (2) the Proprietary Internal Emails (PIE) collection is a proprietary dataset of Microsoft internal emails obtained through an employee participation program. We perform cross-validation on the collection level. That is, when testing on one collection, models are trained and hyperparameters are selected on the other collection (i.e., train/validate on Avocado, test on PIE and vice versa). Models should generalize over multiple tenants (i.e., organizations) as maintaining specialized models is cumbersome. In addition, model effectiveness should remain constant over time to avoid frequent model retraining. Consequently, topical regularities contained within a tenant should not influence our comparison. Furthermore, privacy concerns

may dictate that training and test tenants are different. On the training set, we create a temporal 95/5 split for training and validation/model selection. On the test set, all instances are used for testing only. Attachable items consist of file attachments and URLs; see Table 4.2.

The training and test instances are extracted, for every collection independently, in the following unsupervised manner. File attachments and normalized URLs are extracted from all messages. We remove outlier items by trimming the bottom and top 5% of the attachable item frequency distribution. Infrequent items are non-retrievable and are removed in accordance to our experimental design (Section 4.5.2). However, in this chapter we are interested in measuring the performance on retrieving attachable items that are in the “torso” of the distribution and, consequently, frequent items (e.g., electronic business cards) are removed as well. Any message that links to an attachable item (i.e., URL or attachment) and the message preceding it is considered a request/reply instance. In addition, we filter request/reply instances containing attachable items that (a) occurred previously in the same thread, or (b) contain attachable items that do not occur in the user’s mailbox before the time of the request message (see Section 4.5.2). Mailboxes are indexed and searched using Indri [238, 262]. For message retrieval, we use the Query-Likelihood Model (QLM) with Dirichlet smoothing [286] where the smoothing parameter (μ) is set to the average message length [20, 272]. At test time, query formulation methods extract query terms from the request message, queries are executed using the email search engine of the user (i.e., Indri) and attachable items are ranked according to Eq. 4.1. Rankings are truncated such that they only contain the top-1000 messages and top-100 attachable items. The ground truth consists of binary relevance labels where items linked in the reply message are relevant.

4.5.4 Methods under comparison

As the attachable item recommendation task is first introduced in this chapter, there exist no methods directly aimed at solving this task. However, as mentioned in the related work section (Section 4.2), there are two areas (prior art search and verbose query reduction) that focus on extracting queries from texts. Consequently, we use computationally tractable methods (Section 4.2) from these areas for comparison: (1) single features, i.e., term frequency (TF), TF-IDF, logTF-IDF, relative entropy (RE), used for prior art retrieval [54, 164, 278, 279] where the top- k unique terms are selected from either the subject, the body or both. Hyperparameters $1 \leq k \leq 15$ and, in the case of RE, $\lambda = 0.1, \dots, 0.9$ are optimized on the validation set, (2) the learning to rank method for query term ranking proposed by Lee et al. [143] for the verbose query reduction task. To adapt this method for our purposes, we use the domain-specific features listed in Table 4.1 (where the representations are obtained by training a Skip-Gram word2vec model with default parameters on the email collection), only consider single-term groups (as higher order term groups are computationally impractical during inference) and use a more-powerful pairwise RankSVM [125] (with default parameters [227]) instead of a pointwise approach. Feature value min-max normalization is performed on the instance-level. The context window width $L = 3, 5, \dots, 15$ is optimized on the validation set. In addition, we consider the following baselines: (3) all terms (Full) are selected from either the subject, the body or both,

(4) random terms, selected from the subject, the body or both, where we either select k unique terms randomly (Random k) or a random percentage p of terms (Random %). Hyperparameters $1 \leq k \leq 15$ and $p = 10\%, 20\%, \dots, 50\%$ are optimized on the validation set. Finally, we consider a pointwise alternative to the CNN model: (5) CNN-p with the logistic function at the output layer (instead of the softmax) and terms are selected if their score exceeds a threshold optimized on the validation set F1 score (instead of the EoR token).

The CNN models are trained for 30 iterations using Adam [133] with $\alpha = 10^{-5}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. The iteration with the lowest data loss on the validation set is selected. Word embeddings are 128-dimensional, the two hidden layers have 512 hidden units each, with dropout ($p = 0.50$) and the softplus activation function. Weights are initialized according to [91]. We set the batch size $|B| = 128$ and regularization lambda $\lambda = 0.1$. The context window width $L = 3, 5, \dots, 15$ is optimized on the validation set. For word embeddings (both as part of the CNN and RankSVM), we consider the top-60k terms. Infrequent terms are represented by a shared representation for the unknown token.

4.5.5 Evaluation measures and significance

To answer **RQ1.4**, we report the Mean Reciprocal Rank (MRR), Normalized Discounted Cumulative Gain (NDCG) and Precision at rank 5 (P@5) evaluation measures computed using `trec_eval`.² For **RQ1.5**, we examine the pairwise differences in terms of Reciprocal Rank (RR). In the case of **RQ1.6**, we measure the relative difference in MRR when removing a feature category. Wherever reported, significance is determined using a two-tailed paired Student t-test.

4.6 Results & discussion

We start by presenting a comparison between methods (**RQ1.4**) on attachable item recommendation, provide an error analysis (**RQ1.5**) and perform a feature importance study (**RQ1.6**) (see Section 4.5.2 for an overview of the experimental design).

4.6.1 Overview of experimental results

RQ1.4 Table 4.3 shows the recommendation results of attachable items in enterprise email collections (Section 4.5.3).

We see that CNN outperforms all other query formulation methods on both enterprise email collections. Significance is achieved (MRR) between CNN and the second-best performing methods: CNN-p and RE (subject), respectively, on the Avocado and PIE. The methods that select terms only from the email subject perform strongly on both collections. Within the set of subject methods (1st part of Table 4.3), we also observe that there is little difference between the methods. In fact, for Avocado, simply taking the subject as query performs better than any of the remaining

²https://github.com/usnistgov/trec_eval

4. Reply With: Proactive Recommendation of Email Attachments

Table 4.3: Comparison of CNN with state-of-the-art query formulation methods (Section 4.5.4) on the Avocado and PIE collections. The numbers reported on Avocado were obtained using models trained/validated on PIE and vice versa (Section 4.5.3). Significance is determined using a paired two-tailed Student t-test (* $p < 0.10$; ** $p < 0.05$) [234] between CNN and the second best performing method (in *italic*).

	Avocado			PIE		
	MRR	NDCG	P@5	MRR	NDCG	P@5
Full field, single features and random (subject)						
Full	0.2286	0.3097	0.0686	0.3338	0.4621	0.1088
TF	0.2280	0.3095	0.0686	0.3315	0.4600	0.1079
TF-IDF	0.2250	0.3073	0.0704	0.3390	0.4663	0.1090
logTF-IDF	0.2280	0.3095	0.0686	0.3315	0.4600	0.1079
RE	0.2223	0.3038	0.0698	<i>0.3391</i>	<i>0.4664</i>	<i>0.1095</i>
Random k	0.2143	0.2932	0.0647	0.3266	0.4553	0.1063
Random %	0.1481	0.2104	0.0467	0.2749	0.4013	0.0889
Full field, single features and random (body)						
Full	0.1248	0.1930	0.0377	0.2115	0.3376	0.0672
TF	0.1025	0.1719	0.0309	0.2094	0.3358	0.0660
TF-IDF	0.1507	0.2213	0.0459	0.2237	0.3481	0.0722
logTF-IDF	0.1109	0.1755	0.0311	0.1914	0.3180	0.0627
RE	0.1441	0.2128	0.0424	0.2198	0.3430	0.0699
Random k	0.0785	0.1394	0.0229	0.1781	0.3078	0.0568
Random %	0.1030	0.1646	0.0325	0.1887	0.3128	0.0606
Full field, single features and random (subject + body)						
Full	0.1995	0.2785	0.0612	0.3087	0.4406	0.0972
TF	0.1783	0.2653	0.0551	0.3005	0.4334	0.0953
TF-IDF	0.2097	0.2933	0.0649	0.3100	0.4397	0.0991
logTF-IDF	0.1858	0.2726	0.0592	0.2747	0.4098	0.0871
RE	0.2138	0.2980	0.0649	0.3200	0.4489	0.1023
Random k	0.1404	0.2148	0.0436	0.2721	0.4076	0.0886
Random %	0.1753	0.2514	0.0520	0.2592	0.3941	0.0822
Learning-to-rank methods (subject + body)						
RankSVM	0.1650	0.2425	0.0497	0.3079	0.4392	0.0980
CNN-p	<i>0.2319</i>	<i>0.3129</i>	<i>0.0708</i>	0.3347	0.4630	0.1087
CNN	0.2455*	0.3313**	0.0770**	0.3492**	0.4744**	0.1123

subject-based methods. Subjects convey relevance and context [274] and can compactly describe the topic of a content request. However, in order to generate better queries, we need to extract additional terms from the email body as email subjects tend to be short.

The same methods that we used to extract terms from the subject perform poorly when only presented with the body of the email (2nd part of Table 4.3). When allowing the methods to select terms from the full email (subject and body), we see a small




Figure 4.5: Query length distribution according to the most prominent methods (**S** and **B** denote subject and body fields, resp.). TF-IDF, RE and RankSVM select a fixed number of terms for all queries, whereas CNNs select a variable number of terms.

increase in retrieval performance (3rd part of Table 4.3) compared to body-only terms. However, none of the methods operating on the full email message manage to outperform the subject by itself.

Our learning to rank (LTR) query term methods (last part of Table 4.3) outperform the subject-based methods (ignoring RankSVM). This comes as little surprise, as the presence of a term in the subject is incorporated as a feature in our models (Table 4.1). The reason why RankSVM, originally introduced for reducing long search queries, performs poorly is due to the fact that its training procedure fails to deal with the long length of emails. That is, RankSVM is trained by measuring the decrease in retrieval effectiveness that occurs from leaving one term out of the full email message (i.e., top-down). This is an error-prone way to score query terms as emails tend to be relatively long (Table 4.2). Conversely, our approach to generate silver query training data (Section 4.4.1) considers groups of query terms and the reference query is constructed in a bottom-up fashion.




Figure 4.6: Per-instance differences in Reciprocal Rank (RR) between the email subject query and the CNN query on Avocado. The plot for PIE is qualitatively similar to the one shown. Positive bars (left) indicate instances where the subject performs better than CNN and vice versa for negative (right).

Fig. 4.5 shows the distribution of generated query lengths for the most prominent methods (TF-IDF, RE, RankSVM and the neural networks). On both collections, subject-oriented methods (TF-IDF, RE) seem to extract queries of nearly all the same length. When considering the full subject field, we see that its length varies greatly with outliers of up to 70 query terms. Methods that extract terms from the email body or the full email generate slightly longer queries than the TF-IDF and RE subject methods. The methods that learn to rank query terms generate the longest queries. While RankSVM selects query terms according to a global rank cut-off, the CNNs select a variable number of terms for every request message. Consequently, it comes as no surprise that we observe high variance within the CNN-generated query lengths. In addition, we observe that CNNs have a similar query length distribution as the subject queries. This is due to the fact that the CNNs actually expand the subject terms, as we will see in the next section.

4.6.2 Analysis of differences

RQ1.5 Fig. 4.6 shows the per-instance differences between CNN and the full email subject as query.

For about 60% of request messages both query generation methods (full subject and CNN) generate queries that perform at the same level. In fact, MRR on this subset of instances is 9% (Avocado) and 17% (PIE) better than the best performance over the full test set (Table 4.3). Do the two methods generate identical queries on this subset? The average Jaccard similarity between the queries extracted by both methods is 0.55 (Avocado) and 0.62 (PIE). This indicates that, while query terms extracted by either method overlap, there is a difference in query terms that does not impact retrieval. Upon examining the differences we find that, for the subject, these terms are stopwords and email subject abbreviations (e.g., RE indicating a reply). In the case of CNN, the difference comes from email body terms that further clarify the request. We find that the CNN builds upon the subject query (excluding stopwords) using terms of the body. However, for the 60% of request instances where no difference in performance is observed (Fig. 4.6), the subject by itself suffices to describe the request.

What can we say about the remaining 40% of queries where there is an observable




Figure 4.7: Feature ablation study for the CNN model on the Avocado (stripes) and PIE (dots) benchmarks. One of the following feature categories (Table 4.1) is systematically left out: Message (M), Collection statistics (C), Part-of-Speech (PoS), t_{term} (t), context (c) or all representations (t + c).

difference? A closer look at Fig. 4.6 shows that there are extreme peaks at both sides of the graph and that neither method fully dominates the other. Upon examining the outliers, we find the following trends: (1) When the subject is indescriptive of the email content (e.g., the subject is “*important issue*”) then the CNN can extract better terms from the email body. (2) Topic drift within conversations negatively impacts the retrieval effectiveness of the subject as a query. However, long threads do not necessarily exhibit topic drift as in some cases the subject remains a topical representation of the conversation. (3) Mentions of named entities constitute effective query terms. For example, the name of a person who is responsible for an attachable item tends to improve retrieval effectiveness over using the subject as a query, (4) Long queries generated by the CNN can often disambiguate a request and perform much better than the subject query. (5) In most cases where the subject query outperforms the CNN, this is due to the fact that the CNN model extracts too many noisy terms and creates query drift.

4.6.3 Feature importance

RQ1.6 Fig. 4.7 depicts a feature ablation study where we systematically leave out one of the feature categories.

We observe that both local (message and part-of-speech) and global (collection) features are of importance. When comparing the behavior of the enterprise email collections (Section 4.5.3), we see that the message (M) features have a significant ($p < 0.10$) impact on both collections. The collection statistics (C) yield a significant ($p < 0.10$) difference in the case of PIE; no significant differences were observed in the remaining cases. In addition, while Avocado benefits from the learned representations, the inclusion of the representations of the context slightly decreases performance on PIE. This can be explained by our evaluation setup (Section 4.5.2) where models evaluated on PIE are trained using Avocado and vice versa. Therefore, it is likely that the model learns certain patterns present from the data-scarce Avocado collection (Table 4.2) that causes false positive terms to be selected for PIE.

4.7 Summary

We introduced a novel proactive retrieval task for recommending email attachments that involves formulating a query from an email request message. An evaluation framework was proposed that extracts labeled request/attachment instances from an email corpus containing request/reply pairs automatically. Candidate queries, which we refer to as *silver queries*, are algorithmically synthesized for request/attachment instances and a deep convolutional neural network (CNN) is trained using the silver queries that learns to extract query terms from request messages.

We find that our framework extracts instances that are usable for training and testing. Our CNN, which we train using silver queries, significantly outperforms existing methods for extracting query terms from verbose texts. Terms occurring in the subject of the email are representative of the request and formulating a query using the subject is a strong baseline. A study of the per-instance MRR differences show that the CNN and subject query perform quite differently for 40% of instances. A qualitative analysis suggests that our CNN outperforms the subject query in cases where the subject is indescriptive. In addition, mentions of named entities constitute good query terms and lengthy queries disambiguate the request. In cases when the subject query outperforms the CNN, it is due to noisy terms being selected from the email body. A feature ablation study shows that both local (i.e., message) and global (i.e., collection) features are important.

Our work has the following limitations. (1) In this chapter we only consider terms occurring in the request message as candidates. While this prevents the term candidate set from becoming too large, it does limit the ability for methods to formulate expressive queries in the case where request messages are concise. (2) The retrieval model used in this chapter, a language model with Dirichlet smoothing, is ubiquitous in retrieval systems. However, smoothing allows the search engine to deal with verbose queries [286] that contain terms absent from the messages. Subjects often contain superfluous terms (e.g., email clients prepend *FW* to the subjects of forwarded messages). Consequently, our findings may change when considering other retrieval model classes, such as boolean models or semantic matching models.

Part II

Latent Vector Spaces

5

Unsupervised, Efficient and Semantic Expertise Retrieval

5.1 Introduction

The transition to the knowledge and information economy [1] introduces a great reliance on cognitive capabilities [203]. It is crucial for employers to facilitate information exchange and to stimulate collaboration [68]. In the past, organizations would set-up special-purpose database systems for their members to maintain a profile [25]. However, these systems required employees to be proactive. In addition, self-assessments are known to diverge from reality [36, 138] and document collections can quickly become practically infeasible to manage manually. Therefore, there has been an active interest in automated approaches for constructing expertise profiles [25, 245] and retrieving experts from an organization's heterogeneous document repository [63]. *Expert finding* (also known as expertise retrieval or expert search) addresses the task of finding the right person with the appropriate skills and knowledge [23]. It attempts to provide an answer to the question:

Given a topic X , who are the candidates with the most expertise w.r.t. X ?

The expertise retrieval task gained popularity in the research community during the TREC Enterprise Track [245] and has remained relevant ever since, while broadening to social media and to tracking the dynamics of expertise [21, 23, 36, 71, 82–84, 187, 200, 251]. Existing methods fail to address key challenges: (1) Queries and expert documents use different representations to describe the same concepts [115, 146]. Term mismatches between queries and experts [146] occur due to the inability of widely used maximum-likelihood language models to make use of *semantic* similarities between words [223]. (2) As the amount of available data increases, the need for more powerful approaches with greater learning capabilities than smoothed maximum-likelihood language models is obvious [266]. (3) Supervised methods for expertise retrieval [84, 187] were introduced at the turn of the last decade. However, the acceleration of data availability has the major disadvantage that, in the case of supervised methods, manual annotation efforts need to sustain a similar order of growth. This calls for the further development of *unsupervised* methods. (4) In some expertise retrieval methods, a language model is constructed for every document in the collection. These methods

lack *efficient* query capabilities for large document collections, as each query term needs to be matched against every document [23]. Our proposed solution has a strong emphasis on *unsupervised model construction*, *efficient query capabilities* and *semantic matching* between query terms and candidate experts.

Specifically, we propose an unsupervised log-linear model with efficient inference capabilities for the expertise retrieval task. We show that our approach improves retrieval performance compared to vector space-based and generative language models, mainly due to its ability to perform semantic matching [146]. Our method does not require supervised relevance judgements and is able to learn from raw textual evidence and document-candidate associations alone. The purpose of this chapter is to provide insight in how discriminative language models can improve performance of core retrieval tasks compared to maximum-likelihood language models. Therefore, we avoid explicit feature engineering and the incorporation of external evidence in this chapter. In terms of performance, the current best-performing formal language model [20] exhibits a worst-case time complexity linear in the size of the document collection. In contrast, the inference time complexity of our approach is asymptotically bounded by the number of candidate experts.

The research questions we ask in this chapter towards answering **RQ2** are as follows:

RQ2.1 How does our discriminative log-linear model compare to vector space-based methods and generative language models for the expert retrieval task in terms of retrieval performance?

RQ2.2 What can we learn regarding the different types of errors made by generative and discriminative language models?

RQ2.3 How does the complexity of inference in our log-linear model compare to vector-space based and generative models?

RQ2.4 How does the log-linear model handle incremental indexing and what are its limitations?

5.2 Related work

We refer to Section 2.2 of our background chapter (Chapter 2). Particularly relevant to this chapter is the subsection that covers prior work on expert retrieval (Section 2.2.2.1) and its relation to document retrieval, followed by semantic matching methods (Section 2.2.1) and neural language models (Section 2.2.3).

What we add on top of the related work described above is the following. In this chapter we model the conditional probability of the expertise of a candidate given a single query term (contrary to binary relevance given a character-based n-gram [118]). In the process we learn a distributed vector representation (similar to LSI, pLSI and semantic hashing) for both words and candidates such that nearby representations indicate semantically similar concepts.

We propose a log-linear model that is similar to neural language models. The important difference is that we predict a candidate expert instead of the next word. To the best of our knowledge, we are the first to propose such a solution. We employ an

embedding layer in our shallow model for the same reasons as mentioned above: we learn continuous word representations that incorporate semantic and syntactic similarity tailored to an expert's domain.

5.3 A log-linear model for expert search

In the setting of this chapter we have a document collection D and a predefined set of candidate experts C (entities to be retrieved). Documents $d \in D$ are represented as a sequence of words $w_1, \dots, w_{|d|}$ originating from a vocabulary V , where $w_i \in V$ and the operator $|\cdot|$ denotes the document length in tokens. For every document $d \in D$ we write C_d to denote the set of candidates $c \in C$ associated with it (i.e., $C = \bigcup_{d \in D} C_d$). These document-candidate associations can be obtained explicitly from document meta-data (e.g., the author of an e-mail) or implicitly by mining references to candidates from the document text. Notice that some documents might not be associated with any candidate. When presented with a query q of constituent terms $t_1, \dots, t_{|q|}$, the expert retrieval task is to return a list of candidates $\rho(C)$ ordered according to topical expertise. We generate this ranking using a relatively shallow neural network which directly models $P(\mathbf{c} \mid q)$.

We employ vector-based distributed representations [115], for both words (i.e., word embeddings) and candidate experts, in a way that motivates the unsupervised construction of features that express regularities of the expertise finding domain. These representations can capture the similarity between concepts (e.g., words and candidate experts) by the closeness of their representations in vector space. That is, concepts with similar feature activations are interpreted by the model as being similar, or even interchangeable.

5.3.1 The model

To address the expert search task, we model $P(c_j \mid q)$ and rank candidates c_j accordingly for a given q . We propose an unsupervised, discriminative approach to obtain these probabilities. We construct our model solely from textual evidence: we do not require query-candidate relevance assessments for training and do not consider external evidence about the corpus (e.g., different weightings for different sub-collections), the document (e.g., considering certain parts of the document more useful) nor link-based features.

Let e denote the size of the vector-based distributed representations of both words in V and candidate experts in C . These representations will be learned by the model using gradient descent [177] (Section 5.3.2). For notational convenience, we write $P(\mathbf{c} \mid \cdot)$ for the (conditional) probability distribution over candidates, which is the result of vector arithmetic. We define the probability of a candidate c_j given a single word $w_i \in V$ as the log-linear model

$$P(\mathbf{c} \mid w_i) = \frac{1}{Z_1} \exp(W_c \cdot (W_p \cdot \mathbf{v}_i) + \mathbf{b}_c), \quad (5.1)$$

where W_p is the $e \times |V|$ projection matrix that maps the one-hot representation (i.e., $1\text{-of-}|V|$) of word w_i , \mathbf{v}_i , to its e -dimensional distributed representation, \mathbf{b}_c is a $|C|$ -dimensional bias vector and W_c is the $|C| \times e$ matrix that maps the word embedding

to an unnormalized distribution over candidates C , which is then normalized by $Z_1 = \sum_{j=1}^{|C|} [\exp(W_c \cdot (W_p \cdot \mathbf{v}_i) + \mathbf{b}_c)]_j$. If we consider Bayes' theorem, the transformation matrix W_c and bias vector \mathbf{b}_c can be interpreted as the term log-likelihood $\log P(w_i | \mathbf{c})$ and candidate log-prior $\log P(\mathbf{c})$, respectively. The projection matrix W_p attempts to soften the curse of dimensionality introduced by large vocabularies V and maps words to word feature vectors [32]. Support for large vocabularies is crucial for retrieval tasks [118, 223].

We then assume conditional independence of a candidate's expertise given an observation of data (i.e., a word). Given a sequence of words w_1, \dots, w_k we have:

$$\begin{aligned} P(\mathbf{c} | w_1, \dots, w_k) &= \frac{1}{Z_2} \tilde{P}(\mathbf{c} | w_1, \dots, w_k) = \frac{1}{Z_2} \prod_{i=1}^k P(\mathbf{c} | w_i) \\ &= \frac{1}{Z_2} \exp \left(\sum_{i=1}^k \log (P(\mathbf{c} | w_i)) \right) \end{aligned} \quad (5.2)$$

where $\tilde{P}(\mathbf{c} | w_1, \dots, w_k)$ denotes the unnormalized score and

$$Z_2 = \sum_{j=1}^{|C|} \exp \left(\sum_{i=1}^k \log (P(c_j | w_i)) \right)$$

is a normalizing term. The transformation to log-space in (5.2) is a well-known trick to prevent floating point underflow [185, p. 445]. Given (5.2), inference is straightforward. That is, given query $q = t_1, \dots, t_k$, we compute $P(\mathbf{c} | t_1, \dots, t_k)$ and rank the candidate experts in descending order of probability.

Eq. 5.1 defines a neural network with a single hidden layer. We can add additional layers. Preliminary experiments, however, show that the shallow log-linear model (5.1) performs well-enough in most cases. Only for larger data sets did we notice a marginal gain from adding an additional layer between projection matrix W_p and the softmax layer over C (W_c and bias \mathbf{b}_c), at the expense of longer training times and loss of transparency.

5.3.2 Parameter estimation

The matrices W_p , W_c and the vector \mathbf{b}_c in (5.1) constitute the parameters of our model. We estimate them using error back propagation [218] as follows. For every document $d_j \in D$ we construct an ideal distribution over candidates $\mathbf{p} = P(\mathbf{c} | d_j)$ based on the document-candidate associations C_{d_j} such that

$$P(\mathbf{c} | d_j) = \begin{cases} \frac{1}{|C_{d_j}|}, & c \in C_{d_j} \\ 0, & c \notin C_{d_j} \end{cases}$$

We continue by extracting n-grams where n remains fixed during training. For every n-gram w_1, \dots, w_n generated from document d we compute $\tilde{\mathbf{p}} = P(\mathbf{c} | w_1, \dots, w_n)$ using (5.2). During model constructing we then optimize the cross-entropy $H(\mathbf{p}, \tilde{\mathbf{p}})$

(i.e., the joint probability of the training data if $|C_{d_j}| = 1$ for all j) using batched gradient descent. The loss function for a single batch of m instances with associated targets $(\mathbf{p}^{(i)}, \tilde{\mathbf{p}}^{(i)})$ is as follows:

$$\begin{aligned}
L(W_p, W_c, \mathbf{b}_c) &= \frac{1}{m} \sum_{i=1}^m \frac{|d_{\max}|}{|d^{(i)}|} H(\mathbf{p}^{(i)}, \tilde{\mathbf{p}}^{(i)}) \\
&\quad + \frac{\lambda}{2m} \left(\sum_{i,j} W_{p_{i,j}}^2 + \sum_{i,j} W_{c_{i,j}}^2 \right) \\
&= -\frac{1}{m} \sum_{i=1}^m \frac{|d_{\max}|}{|d^{(i)}|} \sum_{j=1}^{|C|} P(c_j | d^{(i)}) \log \left(P(c_j | w_1^{(i)}, \dots, w_n^{(i)}) \right) \\
&\quad + \frac{\lambda}{2m} \left(\sum_{i,j} W_{p_{i,j}}^2 + \sum_{i,j} W_{c_{i,j}}^2 \right),
\end{aligned} \tag{5.3}$$

where $d^{(i)}$ refers to the document from which n-gram $w_1^{(i)}, \dots, w_n^{(i)}$ was extracted, $d_{\max} = \arg \max_{d \in D} |d|$ indicates the longest document in the collection, and λ is a weight regularization parameter. The update rule for a particular parameter θ (W_p , W_c or \mathbf{b}_c) given a single batch of size m is:

$$\theta^{(t+1)} = \theta^{(t)} - \alpha^{(t)} \odot \frac{\partial L(W_p^{(t)}, W_c^{(t)}, \mathbf{b}_c^{(t)})}{\partial \theta}, \tag{5.4}$$

where $\alpha^{(t)}$ and $\theta^{(t)}$ denote the per-parameter learning rate and parameter θ at time t , respectively. The learning rate α consists of the same number of elements as there are parameters; in the case of a global learning rate, all elements of α are equal to each other. The derivatives of the loss function (5.3) are given in the Appendix.

In the next section we will discuss our experimental setup, followed by an overview of our experimental results and further analysis in Section 5.5.

5.4 Experimental setup

5.4.1 Research questions

As indicated in the introduction of this chapter, we seek to answer the following research questions:

RQ2.1 How does our discriminative log-linear model compare to vector space-based methods and generative language models for the expert retrieval task in terms of retrieval performance?

In particular, how does the model perform when compared against vector space-based (LSI and TF-IDF) and generative approaches (profile-centric Model 1 and document-centric Model 2)?

RQ2.2 What can we learn regarding the different types of errors made by generative and discriminative language models?

Does the best-performing generative model simply perform slightly better on the topics for which the other models perform decent as well, or do they make very different errors? If the latter holds, an ensemble of the rankings produced by both model types might exceed performance of the individual rankings.

RQ2.3 How does the complexity of inference in our log-linear model compare to vector-space based and generative models?

The worst-case inference cost of document-centric models makes them unattractive in online settings where the set of topics is not defined beforehand and the document collection is large. Profile-centric methods are preferred in such settings as they infer from one language model per candidate expert for every topic (i.e., a pseudo-document consisting of a concatenation of all documents associated with an expert) [23]. Vector space-based methods [71] have similar problems due to the curse of dimensionality [122] and consequently their inferential time complexity is likewise asymptotically bounded by the number of experts.

RQ2.4 How does the log-linear model handle incremental indexing and what are its limitations?

5.4.2 Benchmarks

The proposed method is applicable in the setting of the Expert Search task of the TREC Enterprise track from 2005 to 2008 [245]. We therefore evaluate on the W3C and CERC benchmarks released by the track. The W3C dataset [64] is a crawl of the W3C's sites in June 2004 (mailing lists, web pages, etc.). The CSIRO Enterprise Research Collection (CERC) [13] is a dump of the intranet of Australia's national science agency. Additionally, we evaluate our method on a smaller, more recent benchmark based on the employee database of Tilburg University (TU) [36], which consists of bi-lingual, heterogeneous documents. See Table 5.1.

Mining document-candidate associations and how they influence performance has been extensively covered in previous work [20, 23] and is beyond the scope of this chapter. For TU, the associations are part of the benchmark. For W3C, a list of possible candidates is given and we extract the associations ourselves by performing a case-insensitive match of full name or e-mail address [20]. For CERC, we make use of publicly released associations [15].

As evaluation measures we use Mean Average Precision (MAP), Mean Reciprocal Rank (MRR), Normalized Discounted Cumulative Gain at rank 100 (NDCG@100) and Precision at rank 5 (P@5) and rank 10 (P@10).

5.4.3 Baselines

We compare our approach to existing unsupervised methods for expert retrieval that solely rely on textual evidence and static document-candidate associations. (1) Demartini et al. [71] propose a generic framework to adapt vector spaces operating on documents to entities. We compare our method to TF-IDF (raw frequency and inverse

Table 5.1: An overview of the three datasets (W3C, CERC and TU) used for evaluation and analysis.

	W3C	CERC	TU
Number of documents	331,037	370,715	31,209
Average document length ^a	1,237.23	460.48	2,454.93
Number of candidates ^b	715	3,479	977
Number of document-candidate associations	200,939	236,958	36,566
Number of documents (with $C_d > 0$)	93,826	123,934	27,834
Number of associations per document ^c	2.14 ± 3.29	1.91 ± 3.70	1.13 ± 0.39
Number of associations per candidate	281.03 ± 666.63	$68.11 \pm 1,120.74$	37.43 ± 61.00
Queries			
	49 (2005)	50 (2007)	1,662 (GT1)
	50 (2006)	77 (2008)	1,266 (GT5)

^a Measured in number of tokens.^b Only candidates with at least a single document association are considered.^c Only documents with at least one association are considered.

document frequency) and LSI (300 latent topics) variants of their vector space model for entity ranking (using cosine similarity). (2) In terms of language modelling, Balog et al. [20] propose two models for expert finding based on generative language models. The first takes a profile-centric approach comparing the language model of every expert to the query, while the second is document-centric. We consider both models with different smoothing configurations: Jelinek-Mercer (jm) smoothing with $\lambda = 0.5$ [20] and Dirichlet (d) smoothing with β equal to the average document length [21] (see Table 5.1). Significance of results produced by the baselines (compared to our method) is determined using a two-tailed paired randomization test [234].

5.4.4 Implementation details

The vocabulary V is constructed from each corpus by ignoring punctuation, stop words and case; numbers are replaced by a numerical placeholder token. During our experiments we prune V by only retaining the 2^{16} most-frequent words so that each word can be encoded by a 16-bit unsigned integer. Incomplete n-gram instances are padded by a special-purpose token.

In terms of parameter initialization, we sample the initial matrices W_c and W_p (5.1) uniformly in the range

$$\left[-\sqrt{\frac{6.0}{m+n}}, \sqrt{\frac{6.0}{m+n}} \right]$$

for an $m \times n$ matrix, as this initialization scheme improves model training convergence [91], and take the bias vector \mathbf{b}_c to be null. The projection layer W_p is initialized with pre-trained word representations trained on Google News data [176]; the number of word features is set to $e = 300$, similar to pre-trained representations.

We used adadelata ($\rho = 0.95$, $\epsilon = 10^{-6}$) [284] with batched gradient descent ($m = 1024$) and weight decay $\lambda = 0.01$ during training on NVidia GTX480 and NVidia Tesla K20 GPUs. We only iterate once over the entire training set for each experiment.

5.5 Results & discussion

We start by giving a high-level overview of our experimental results and then address issues of scalability, provide an error analysis and discuss the issue of incremental indexing.

5.5.1 Overview of experimental results

We evaluate the log-linear model on the W3C, CERC and TU benchmarks (Section 5.4.2). During training we extract non-overlapping n-grams for the W3C and CERC benchmarks and overlapping n-grams for the TU benchmark. As the TU benchmark is considerably smaller, we opted to use overlapping n-grams to counter data sparsity. The architecture of each benchmark model (e.g., number of candidate experts) is inherently specified by the benchmarks themselves (see Table 5.1). However, the choice of n-gram size during training remains open. Errors for input w_1, \dots, w_n are

Table 5.2: Evaluation results for models trained on the W3C, CERC and TU benchmarks. Suffixes (d) and (jm) denote Dirichlet and Jelinek-Mercer smoothing, respectively (Section 5.4.3). Significance of results is determined using a two-tailed paired randomization test [234] (***) $p < 0.01$; ** $p < 0.05$; * $p < 0.1$ with respect to the log-linear model (adjusted using the Benjamini-Hochberg procedure for multiple testing [34]).

	W3C	2005					2006				
		MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR	P@5	P@10
LSI		0.135	0.266	0.306	0.192	0.196	0.245	0.371	0.482	0.287	0.338
TF-IDF		0.243	0.426	0.541	0.384	0.350	0.343	0.531	0.650	0.492	0.498
Model 1 (d)		0.192	0.358	0.433	0.276	0.266	0.321	0.491	0.635	0.478	0.449
Model 1 (jm)		0.190	0.352	0.390	0.272	0.276	0.311	0.483	0.596	0.502	0.437
Model 2 (d)		0.198	0.369	0.429	0.288	0.272	0.261	0.419	0.551	0.441	0.404
Model 2 (jm)		0.211	0.380	0.451	0.332	0.296	0.260	0.423	0.599	0.449	0.429
Log-linear (ours)		0.248	0.444	0.618*	0.412	0.361	0.484**	0.667***	0.833***	0.713***	0.644***
	CERC	2007					2008				
		MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR	P@5	P@10
LSI		0.031	0.107	0.060	0.016	0.014	0.038	0.099	0.106	0.042	0.055
TF-IDF		0.332	0.486	0.463	0.196	0.141	0.269	0.465	0.525	0.332	0.277
Model 1 (d)		0.287	0.427	0.384	0.156	0.096	0.181	0.355	0.388	0.200	0.172
Model 1 (jm)		0.278	0.420	0.384	0.156	0.084	0.170	0.347	0.339	0.181	0.159
Model 2 (d)		0.352	0.495	0.454	0.180	0.138	0.264	0.461	0.510	0.281	0.244
Model 2 (jm)		0.361	0.500	0.467	0.192	0.138	0.274	0.463	0.517	0.278	0.239
Log-linear (ours)		0.344	0.493	0.513	0.215	0.150	0.342***	0.519**	0.656**	0.381*	0.299
	TU	GT1					GT5				
		MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR	P@5	P@10
LSI		0.095	0.205	0.153	0.060	0.051	0.097	0.208	0.129	0.043	0.036
TF-IDF		0.216	0.356	0.324	0.131	0.097	0.233	0.378	0.288	0.108	0.079
Model 1 (d)		0.171	0.308	0.258	0.103	0.082	0.241	0.385	0.292	0.109	0.081
Model 1 (jm)		0.189	0.325	0.277	0.112	0.085	0.231	0.373	0.271	0.100	0.075
Model 2 (d)		0.154	0.284	0.228	0.087	0.070	0.191	0.334	0.233	0.084	0.065
Model 2 (jm)		0.234	0.370	0.342	0.136	0.101	0.253	0.394	0.302	0.108	0.081
Log-linear (ours)		0.219	0.356	0.351	0.145*	0.105	0.287***	0.425***	0.363***	0.134***	0.092***

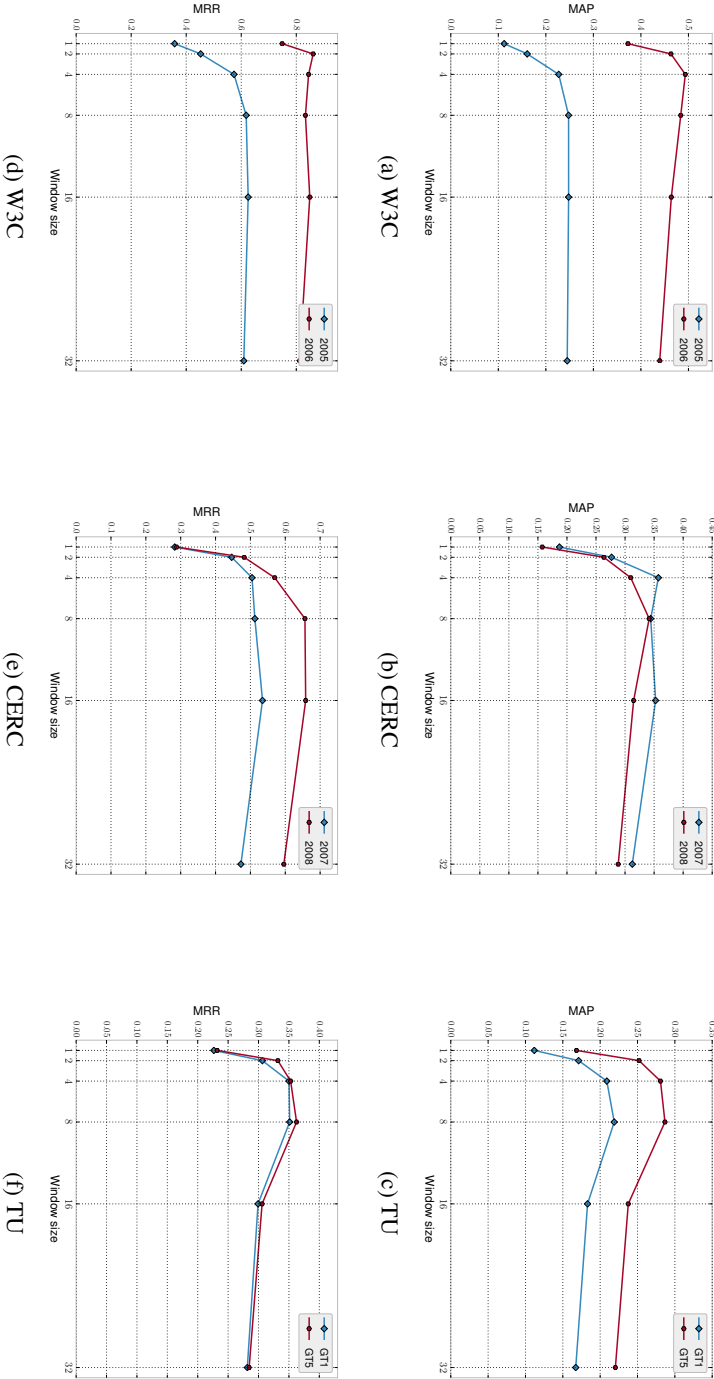


Figure 5.1 : Sensitivity analysis for window size (n-gram) during parameter estimation (5.3) for W3C, CERC and TU benchmarks.

propagated back through W_c until the projection matrix W_p is reached; if a single word w_i causes a large prediction error, then this will influence its neighbouring words $w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_n$ as well. This allows the model to learn continuous word representations tailored to the expert retrieval task and the benchmark domain.

A larger window size has a negative impact on batch throughput during training. We are thus presented with the classic trade-off between model performance and construction efficiency. Notice, however, that the number of n-grams decreases as the window size increases if we extract non-overlapping instances. Therefore, larger values of n lead to faster wall-clock time model construction for the W3C and CERC benchmarks in our experiments.

We sweep over the window width $n = 2^i$ ($0 \leq i < 6$) for all three benchmarks and their corresponding relevance assessments. We report MAP and MRR for every configuration (see Fig. 5.1). We observe a significant performance increase between $n = 1$ and $n = 2$ on all benchmarks, which underlines the importance of the window size parameter. The increase in MAP implies that the performance achieved is not solely due to initialization with pre-trained representations (Section 5.4.4), but that the model efficiently learns word representations tailored to the problem domain. The highest MAP scores are attained for relatively low n . As n increases beyond $n = 8$ a gradual decrease in MAP is observed on all benchmarks. In our remaining experiments we choose $n = 8$ regardless of the benchmark.

Words w_i that mainly occur in documents associated with a particular expert are learned to produce distributions $P(c | w_i)$ with less uncertainty than words associated with many experts in (5.1). The product of $P(c | w_i)$ in (5.2) aggregates this expertise evidence generated by query terms. Hence, queries with strong evidence for a particular expert should be more predictable than very generic queries. To quantify uncertainty we measure the normalized entropy [230] of $P(c | q)$:

$$\eta(c | q) = -\frac{1}{\log(|C|)} \sum_{j=1}^{|C|} p(c_j | q) \log(p(c_j | q)). \quad (5.5)$$

Equation 5.5 can be interpreted as a similarity measure between the given distribution and the uniform distribution. Importantly, Fig. 5.2 shows that there is a statistically significant negative correlation between query-wise normalized entropy and average precision for all benchmarks.

Table 5.2 presents a comparison between the log-linear model and the various baselines (Section 5.4.3). Our unsupervised method significantly ($p < 0.01$) outperforms the LSI-based method consistently. In the case of the TF-IDF method and the profile-centric generative language models (Model 1), we always perform better and statistical significance is achieved in the majority of cases. The document-centric language models (Model 2) perform slightly better than our method on two (out of six) benchmarks in terms of MAP and NDCG@100: (a) For the CERC 2007 assessment we match performance of the document-centric generative model with Jelinek-Mercer smoothing. (b) For TU GT1 the generative counterpart seems to outperform our method.

Notice that over all assessments, the log-linear model consistently outperforms all profile-centric approaches and is only challenged by the smoothed document-centric




Figure 5.2: Scatter plot of the normalized entropy of distribution $P(c | q)$ (5.2) returned by the log-linear model and per-query average precision for W3C, CERC and TU benchmarks. Pearson's R and associated p -value (two-tailed paired permutation test: $*** p < 0.01$, $** p < 0.05$, $* p < 0.1$) are between parentheses. The depicted linear fit was obtained using an ordinary least squares regression.

approach. In addition, for the precision-based measures ($P@k$ and MRR), the log-linear model consistently outperforms all other methods we compare to.

Next, we turn to a topic-wise comparative analysis of discriminative and generative models. After that, we analyse the scalability and efficiency of the log-linear model and compare it to that of the generative counterparts, and address incremental indexing.

5.5.2 Error analysis

How does our log-linear model achieve its superior performance over established generative models? Fig. 5.3 depicts the per-topic differences in average precision between the log-linear model and Model 2 (with Jelinek-Mercer smoothing) on all benchmarks. For each plot, the vertical bars with a positive AP difference correspond to test topics for which the log-linear model outperforms Model 2 and vice versa for bars with a negative AP difference.

The benefit gained from the projection matrix W_p is two-fold. First, it avoids the curse of dimensionality introduced by large vocabularies. Second, term similarity with respect to the expertise domain is encoded in latent word features. When examining words nearby query terms in the embedding space, we found words to be related to the query term. For example, word vector representations of *xml* and *nonterminal* are very similar for the W3C benchmark (l_2 norm). This can be further observed in Fig. 5.1: log-linear models trained on single words perform significantly worse compared to those that are able to learn from neighbouring words.

We now take a closer look at the topics for which the log-linear model outperforms Model 2 and vice versa. More specifically, we investigate textual evidence related to a topic and whether it is considered relevant by the benchmark. For the log-linear model, we examine terms nearby topic terms in W_p (l_2 -norm), as these terms are considered semantically similar by the model and provide a means for semantic matching. For every benchmark, we first consider topics where exact matches (Model 2) perform best, followed by examples which benefit from semantic matching (log-linear model). Topic identifiers are between parentheses.

W3C Topics *P3P specification* and *CSS3* (EX8 and EX69, respectively) should return candidates associated with the definition of these standards. The log-linear model, however, considers these close to related technologies such as *CSS2* for *CSS3* and *UTF-8* for *P3P*. Semantic matching works for topics *Semantic Web Coordination* and *Annotea server protocol* (EX1 and EX103), where the former is associated with RDF libraries, RDF-related jargon and the names of researchers in the field, while the latter is associated with implementations of the protocol and the maintainer of the project.

CERC For CSIRO, topic *nanohouse* (CE-035) is mentioned in many irrelevant contexts (i.e., spam) and therefore semantic matching fails. The term *fish oil* (CE-126) is quickly associated with different kinds of fish, oils and organizations related to marines and fisheries. On the other hand, we observe *sensor networks* (CE-018) to be associated with sensor/networking jargon and sensor platforms. Topic *forensic science workshop* (CE-103) expands to syntactically-similar terms (e.g., plural), the names of science laboratories and references to support/law-protection organizations.

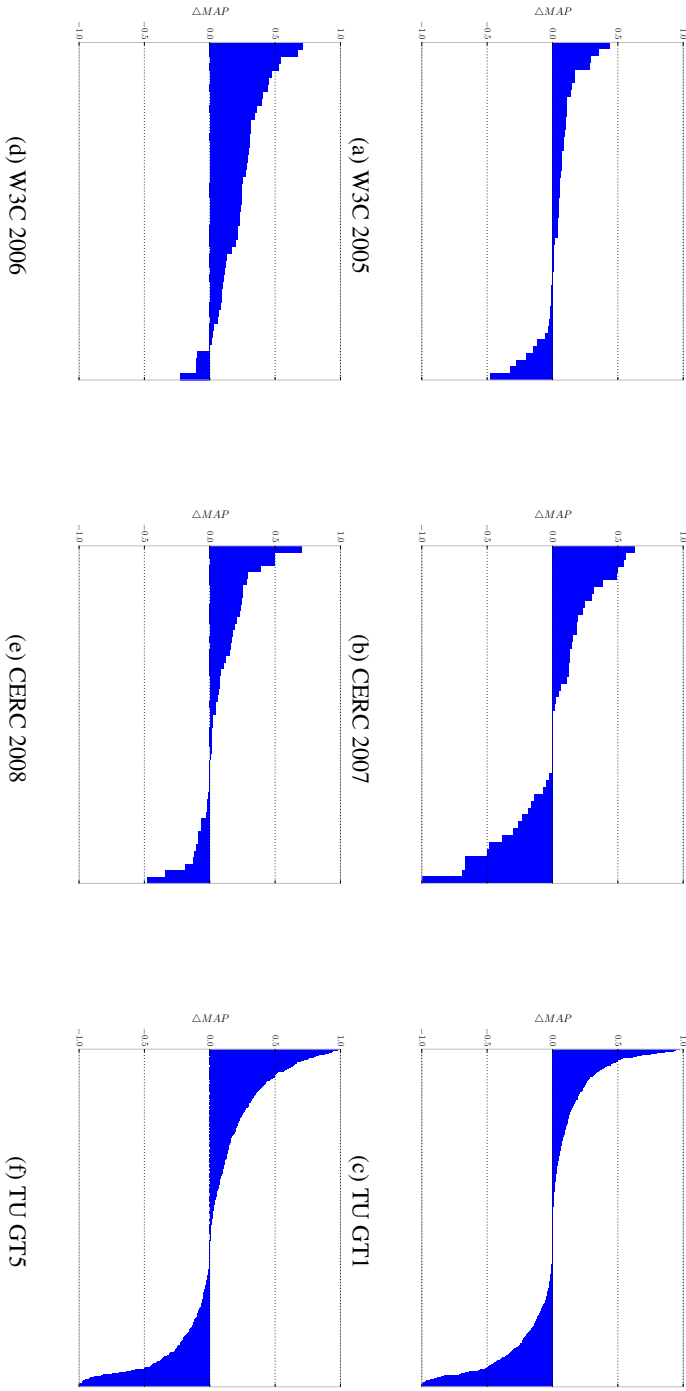


Figure 5.3: Difference of average precision between log-linear model and Model 2 [20] with Jelinek-Mercer smoothing per topic for W3C, CERC and TU benchmarks.

TU The TU benchmark contains both English and Dutch textual evidence. Topics *sustainable tourism* and *interpolation* (1411 and 4882) do not benefit from semantic matching due to a semantic gap: *interpolation* is associated with the polynomial kind while the relevance assessments focus on stochastic methods. Interestingly, for the topic *law and informatization/computerization* (1719) we see that the Dutch translation of *law* is very closely related. Similar terms to *informatization* are, according to the log-linear model, Dutch words related to cryptography. Similar dynamics are at work for *legal-political space* (12603), where translated terms and semantic-syntactic relations aid performance.

In order to further quantify the effect of the embedding matrix W_p , we artificially expand benchmark topic terms by k nearby terms. We then examine how the performance of a profile-centric generative language model [20, Model 1] evolves for different values of k (Fig. 5.4). The purpose of this analysis is to provide further insight in the differences between maximum-likelihood language models and the log-linear model. Fig. 5.4 shows that, for most benchmarks, MAP increases as k goes up. Interestingly enough, the two benchmarks that exhibit a decrease in MAP for larger k (CERC 2007 and TU GT1) are likewise those for which generative language models outperform the log-linear model in Table 5.2. This suggests that the CERC 2007 and TU GT1 benchmarks require exact term matching, while the remaining four benchmarks benefit greatly from the semantic matching provided by our model.

The per-topic differences suggest that Model 2 and the log-linear model make very different errors: Model 2 excels at retrieving exact query matches, while the log-linear model performs semantic matching. Based on these observations we hypothesize that a combination of the two approaches will raise retrieval performance even further. To test this hypothesis, we propose a simple ensemble of rankings generated by Model 2 and the log-linear model by re-ranking candidates according to the multiplicatively-combined reciprocal rank:

$$\text{rank}_{\text{ensemble}}(c_j, q_i) \propto \frac{1}{\text{rank}_{\text{model 2}}(c_j, q_i)} \cdot \frac{1}{\text{rank}_{\text{log-linear}}(c_j, q_i)}, \quad (5.6)$$

where $\text{rank}_M(c_j, q_i)$ denotes the position of candidate c_j in a ranking generated by model M for answering query q_i . Equation (5.6) is equivalent to performing data fusion using CombSUM [231] where the scores are given by the logarithm of the reciprocal ranks of the experts. Table 5.3 compares the result of this ensemble to that of its constituents. Compared to the supervised methods of Fang et al. [84], we conclude that our fully unsupervised ensemble matches the performance of their method on the CERC 2007 benchmark and outperforms their method on the W3C 2005 benchmark. The superior performance of the ensemble suggests the viability of hybrid methods that combine semantic and exact matching.

5.5.3 Scalability and efficiency

Inference in the log-linear model is expressed in linear algebra operations (Section 5.3). These operations can be efficiently performed by highly optimized software libraries and special-purpose hardware (i.e., GPUs). But the baseline methods against which




Figure 5.4: Effect of query expansion by adding nearby terms in W_p (5.1) in traditional language models (Model 1 [20] with Jelinek-Mercer smoothing) for W3C, CERC and TU benchmarks.

Table 5.3: Comparison of Model 2, the log-linear model and an ensemble of the former on W3C, CERC and TU benchmarks. Significance of results is determined using a two-tailed paired randomization test [234] (** $p < 0.01$; *** $p < 0.05$; * $p < 0.1$) with respect to the ensemble ranking (adjusted using the Benjamini-Hochberg procedure for multiple testing [34]).

W3C	2005				2006			
	MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR
Model 2 (jm)	0.211	0.380	0.451	0.332	0.296	0.260	0.423	0.599
Log-linear (ours)	0.248	0.444	0.618	0.412	0.361	0.484 ***	0.667 **	0.833
Ensemble	0.291 ***	0.479 **	0.668	0.440	0.378	0.433	0.634	0.825
CERC								
	2007				2008			
	MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR
Model 2 (jm)	0.361	0.500	0.467	0.192	0.138	0.274	0.463	0.517
Log-linear (ours)	0.344	0.493	0.513	0.215	0.150	0.342	0.519	0.656
Ensemble	0.452 **	0.589 ***	0.627 ***	0.248 *	0.160	0.395 ***	0.593 ***	0.716
TU								
	GT1				GT5			
	MAP	NDCG@100	MRR	P@5	P@10	MAP	NDCG@100	MRR
Model 2 (jm)	0.234	0.370	0.342	0.136	0.101	0.253	0.394	0.302
Log-linear (ours)	0.219	0.356	0.351	0.145	0.105	0.287	0.425	0.363
Ensemble	0.271 ***	0.417 ***	0.403 ***	0.165 ***	0.121 ***	0.331 ***	0.477 ***	0.402 ***
								0.156
								0.105 ***

we compare do not benefit from these speed-ups. Furthermore, many implementation-specific details and choice of parameter values can influence runtime considerably (e.g. size of the latent representations). Therefore, we opt for a theoretical comparison of the inference complexity of the log-linear model and compare these to the baselines (Section 5.4.3).

The log-linear model generates a ranking of candidate experts by straight-forward matrix operations. The look-up operation in the projection matrix W_p occurs in constant time complexity, as the multiplication with the one-hot vector v_i comes down to selecting the i -th column from W_p . Multiplication of the $|C| \times e$ matrix W_c with the e -dimensional word feature vector exhibits $O(|C| \cdot e)$ runtime complexity. If we consider addition of the bias term and division by the normalizing function Z_1 , the time complexity of (5.1) becomes

$$O(\underbrace{|C| \cdot (e + (e - 1))}_{\text{matrix-vector multiplication}} + \underbrace{|C|}_{\text{bias term}} + \underbrace{2 \cdot |C| - 1}_{Z_1}).$$

Notice, however, that the above analysis considers sequential execution. Modern computing hardware has the ability to parallelize common matrix operations [85, 139]. The number of candidate experts $|C|$ is the term that impacts performance most in the log-linear model (under the assumption that $|C| \gg e$).

If we consider n terms, where n is the query length during inference or the window size during training, then the complexity of (5.2) becomes

$$\underbrace{O(n \cdot |C| \cdot (2 \cdot e - 1) + n \cdot (3 \cdot |C| - 1))}_{n \text{ forward-passes}} + \underbrace{(n - 1) \cdot |C|}_{\text{factor product}} + \underbrace{2 \cdot |C| - 1}_{Z_2}$$

Notice that Z_2 does not need to be computed during inference as it does not affect the candidate expert ranking.

In terms of space complexity, parameters W_p , W_c and \mathbf{b}_e , in addition to the intermediate results, all require memory space proportional to their size. Considering (5.2) for a sequence of k words and batches of m instances, we require $O(m \cdot k \cdot |C|)$ floating point numbers for every forward-pass to fit in-memory. While such an upper bound seems reasonable by modern computing standards, it is a severely limiting factor when considering large-scale communities and while utilizing limited-memory GPUs for fast computation.

The inferential complexity of the vector space-based models for entity retrieval [71] depends mainly on the dimensionality of the vectors and the number of candidate experts. The dimensionality of the latent entity representations is too high for efficient nearest neighbour retrieval [122] due to the curse of dimensionality. Therefore, the time complexity for the LSI- and TF-IDF-based vector space models are respectively $O(\gamma \cdot |C|)$ and $O(|V| \cdot |C|)$, where γ denotes the number of latent topics in the LSI-based method. As hyperparameters e and γ both indicate the dimensionality of latent entity representations, the time complexity of the LSI-based method is comparable to that of the log-linear model. We note that $|V| \gg |C|$ for all benchmarks ($|V|$ is between 18 to 91 times larger than $|C|$) we consider in this section and therefore conclude that the TF-IDF method loses to the log-linear model in terms of efficiency.

Compared to the unsupervised generative models of Balog et al., we have the profile-centric Model 1 and the document-centric Model 2 with inference time complexity $O(n \cdot |C|)$ and $O(n \cdot |D|)$, respectively, with $|D| \gg |C|$. In the previous section we showed that the log-linear model always performs better than Model 1 and nearly always outperforms Model 2. Hence, our log-linear model generally achieves the expertise retrieval performance of Model 2 (or higher) at the complexity cost of Model 1 during inference.

5.5.4 Incremental indexing

Existing unsupervised methods use well-understood maximum-likelihood language models that support incremental indexing. We now briefly discuss the incremental indexing capabilities of our proposed method. Extending the set of candidate experts C requires the log-linear model to be re-trained from scratch as it changes the topology of the network. Moreover, every document associated with a candidate expert is considered as a negative example for all other candidates. While it is possible to reiterate over all past documents and only learn an additional row in matrix W_c , the final outcome is unpredictable.

If we consider a stream of documents instead of a predefined set D , the log-linear model can be learned in an online fashion. However, stochastic gradient descent requires that training examples are picked at random such that the batched update rule (5.4) behaves like the empirical expectation over the full training set [42]. While we might be able to justify the assumption that documents arrive randomly, the n -grams extracted from those documents clearly violate this requirement.

Considering a stream of documents leads to the model forgetting expertise evidence as an (artificial) shift in the underlying distribution of the training data occurs. While such behaviour is undesirable for the task considered in this section, it might be well-suited for temporal expert finding [83, 221], where expertise drift over time is considered. However, temporal expertise finding is beyond the scope for this section and left for future work.

5.6 Summary

We have introduced an unsupervised discriminative, log-linear model for the expert retrieval task. Our approach exclusively employs raw textual evidence. Future work can focus on improving performance by feature engineering and incorporation of external evidence. Furthermore, no relevance feedback is required during training. This renders the model suitable for a broad range of applications and domains.

We evaluated our model on the W3C, CERC and TU benchmarks and compared it to state-of-the-art vector space-based entity ranking (based on LSI and TF-IDF) and language modelling (profile-centric and document-centric) approaches. The log-linear model combines the ranking performance of the best maximum-likelihood language modelling approach (document-centric) with inference time complexity linear in the number of candidate experts. We observed a notable increase in precision over existing methods. Analysis of our model's output reveals a negative correlation between the

per-query performance and ranking uncertainty: higher confidence (i.e., lower entropy) in the rankings produced by our approach often occurs together with higher rank quality.

An error analysis of the log-linear model and traditional language models shows that the two make very different errors. These errors are mainly due to the semantic gap between query intent and the raw textual evidence. Some benchmarks expect exact query matches, others are helped by our semantic matching. An ensemble of methods employing exact and semantic matching generally outperforms the individual methods. This observation calls for further research in the area of combining exact and semantic matching.

In the next chapter (Chapter 6), we further investigate the representations learned as part of the log-linear model. In particular, we investigate whether (a) a clustering of experts corresponds to the working groups within organizations, and (b) experts that work in the same fields have similar representations. Chapter 7 and Chapter 8 are dedicated to scaling up the latent vector space model introduced in this chapter to retrieval problems that consist of more retrievable objects and with different characteristics than expert finding.

5.7 Appendix

The derivative of (5.3) w.r.t. bias term \mathbf{b}_c equals

$$\frac{\partial L(W_p, W_c, \mathbf{b}_c)}{\partial \mathbf{b}_c} = -\frac{1}{m} \left(\sum_{i=1}^m \frac{|d_{\max}|}{|d^{(i)}|} \sum_{j=1}^{|C|} P(c_j | d^{(i)}) \frac{\partial \log \left(P(c_j | w_1^{(i)}, \dots, w_n^{(i)}) \right)}{\partial \mathbf{b}_c} \right)$$

and w.r.t. an arbitrary matrix parameter θ (W_p or W_c):

$$\begin{aligned} \frac{\partial L(W_p, W_c, \mathbf{b}_c)}{\partial \theta} &= -\frac{1}{m} \left(\sum_{i=1}^m \frac{|d_{\max}|}{|d^{(i)}|} \sum_{j=1}^{|C|} P(c_j | d^{(i)}) \frac{\partial \log \left(P(c_j | w_1^{(i)}, \dots, w_n^{(i)}) \right)}{\partial \theta} \right) \\ &\quad + \frac{\lambda}{m} \sum_{i,j} \theta_{i,j}. \end{aligned}$$

Further differentiation for parameter θ (W_p , W_c or \mathbf{b}_c):

$$\begin{aligned} \frac{\partial \log \left(P(c_j | w_1, \dots, w_n) \right)}{\partial \theta} &= \frac{1}{P(c_j | w_1, \dots, w_n)} \frac{\partial P(c_j | w_1, \dots, w_n)}{\partial \theta} \\ \frac{\partial P(c_j | w_1, \dots, w_n)}{\partial \theta} &= \frac{\frac{\partial \tilde{P}(c_j | w_1, \dots, w_n)}{\partial \theta} Z_2 - \tilde{P}(c_j | w_1, \dots, w_n) \frac{\partial Z_2}{\partial \theta}}{Z_2^2} \\ \frac{\partial Z_2}{\partial \theta} &= \sum_k \frac{\partial \tilde{P}(c_k | w_1, \dots, w_n)}{\partial \theta} \\ \frac{\partial \tilde{P}(c_j | w_1, \dots, w_n)}{\partial \theta} &= \sum_k \frac{\partial P(c_j | w_k)}{\partial \theta} \prod_{i \neq k} P(c_j | w_i) \end{aligned}$$

For a given candidate c_j and word w_i , following (5.1) we have

$$\begin{aligned} P(c_j | w_i) &= \frac{\tilde{P}(c_j | w_i)}{Z_1} \\ &= \frac{\exp((\sum_{k=1}^e W_{c_j,k} W_{p_k,i}) + b_{c_j})}{\sum_{l=1}^{|C|} \exp((\sum_{k=1}^e W_{c_l,k} W_{p_k,i}) + b_{c_l})} \end{aligned}$$

and consequently, with $\mathbf{W}_{\mathbf{p}_i}^\top$ denoting the i -th column of matrix W_p ,

$$\begin{aligned} \frac{\partial P(c_j | w_i)}{\partial \mathbf{W}_{\mathbf{c}_j}} &= \frac{(Z_1 - \tilde{P}(c_j | w_i)) \tilde{P}(c_j | w_i) \mathbf{W}_{\mathbf{p}_i}^\top}{Z_1^2} \\ \frac{\partial P(c_j | w_i)}{\partial b_{c_j}} &= \frac{(Z_1 - \tilde{P}(c_j | w_i)) \tilde{P}(c_j | w_i)}{Z_1^2} \\ \frac{\partial P(c_j | w_i)}{\partial \mathbf{W}_{\mathbf{p}_i}^\top} &= \frac{(\mathbf{W}_{\mathbf{c}_j} - \sum_{l=1}^{|C|} \mathbf{W}_{\mathbf{c}_l}) \tilde{P}(c_j | w_i)}{Z_1} \end{aligned} \tag{5.7}$$

As can be seen in (5.7), the distributed representations of candidates c_j at time $t + 1$ are updated using the representation of words w_i at time t and vice versa.

6

Structural Regularities in Text-based Entity Vector Spaces

6.1 Introduction

The construction of latent entity representations is a recurring problem [39, 57, 71, 114, 289] in natural language processing and information retrieval. So far, entity representations are mostly learned from relations between entities [39, 289] for a particular task in a supervised setting [114]. How can we learn latent entity representations if (i) entities only have relations to documents in contrast to other entities (e.g., scholars are represented by the papers they authored), and (ii) there is a lack of labeled data?

As entities are characterized by documents that consist of words, can we use word embeddings to construct a latent entity representation? Distributed representations of words [115], i.e., word embeddings, are learned as part of a neural language model and have been shown to capture semantic [58] and syntactic regularities [177, 199]. In addition, word embeddings have proven to be useful as feature vectors for natural language processing tasks [248], where they have been shown to outperform representations based on count-based distributional semantics [24]. A down-side of word embeddings [32] is that they do not take into account the document a word sequence occurred in or the entity that generated it.

Le and Mikolov address this problem by extending word2vec models to doc2vec by additionally modelling the document a phrase occurred in. That is, besides word embeddings they learn embeddings for documents as well. We can apply doc2vec to the entity representation problem by representing an entity as a pseudo-document consisting of all documents the entity is associated with. The neural model we introduced in the previous chapter—which we refer to as SERT in this chapter—incorporates real-world structural relations between represented entities even though the representations are learned from text only. In addition to word embeddings, we learned a representation for entities such that the words that are highly discriminative for an entity have a representation similar to that entity.

In this chapter, we study the regularities contained within entity representations that are estimated, in an unsupervised manner, from texts and associations alone. Do they correspond to structural real-world relations between the represented entities? E.g., if the entities we represent are people, do these regularities correspond to collaborative

and hierarchical structures in their domain (industrial, governmental or academic organizations in the case of experts)? Answers to these questions are valuable because if they allow us to better understand the inner workings of entity retrieval models and give important insights into the entity-oriented tasks they are used for [141]. In addition, future work can build upon these insights to extract structure within entity domains given only a document collection and entity-document relations so to complement or support structured information.

Our working hypothesis is that text-based entity representations encode regularities within their domain. To test this hypothesis we compare latent text-based entity representations learned by neural networks (word2vec, doc2vec, SERT), count-based entity vector representations constructed using Latent Semantic Indexing (LSI) and Latent Dirichlet Allocation (LDA), dimensionality-reduced adjacency representations (Graph PCA) and random representations sampled from a standard multivariate normal distribution. For evaluation purposes we focus on expert finding, a particular case of entity ranking. Expert finding is the task of finding the right person with the appropriate skills or knowledge [23], based on a document collection and associations between people and documents. These associations can be extracted using entity linking methods or from document meta-data (e.g., authorship). Typical queries are descriptions of expertise areas, such as *distributed computing*, and expert search engines answer the question “Who are experts on *distributed computing*?” asked by people unfamiliar with the field.

Our main finding is that, indeed, semantic entity representations encode domain regularities. Entity representations can be used as feature vectors for clustering and that those partitions correspond to structural groups within the entity domain. We also find that similarity between entity representations correlates with relations between entities. In particular, we show how representations of experts in the academic domain encode the co-author graph. Lastly, we show that one of the semantic representation learning methods, SERT, additionally encodes importance amongst entities and, more specifically, the hierarchy of scholars in academic institutions.

The research questions we ask in this chapter towards answering **RQ2** are as follows:

RQ2.5 Do clusterings of text-based entity representations reflect the structure of their domains?

RQ2.6 To what extent do different text-based entity representation methods encode relations between entities?

6.2 Related work

We refer to Section 2.2 of our background chapter (Chapter 2). The subsection covering entity retrieval (Section 2.2.2) and regularities in language representations (Section 2.2.3.1) are of special interest to this chapter.

In the maximum-likelihood language modelling paradigm, experts are represented as a normalized bag-of-words vector with additional smoothing. These vectors are high-dimensional and sparse due to the large vocabularies used in expert domains. Therefore, bag-of-words vectors are unsuited for use as representations—the topic of

study in this chapter—as lower-dimensional and continuous vector spaces are preferred in machine learning algorithms [271].

6.3 Text-based entity vector spaces

For text-based entity retrieval tasks we are given a document collection D and a set of entities X . Documents $d \in D$ consist of a sequence of words $w_1, \dots, w_{|d|}$ originating from a vocabulary V , where $|\cdot|$ denotes the document length in number of words. For every document d we have a set $X_d \subset X$ of associated entities (X_d can be empty for some documents) and conversely $D_x \subset D$ consists of all documents associated with entity x . The associations between documents and experts can be obtained in multiple ways. E.g., named-entity recognition can be applied to the documents and mentions can subsequently be linked to entities. Or associations can be extracted from document meta-data (e.g., authorship).

Once determined, the associations between entities X and documents D encode a bipartite graph. If two entities $x_i, x_j \in X$ are associated with the same document, we say that x_i and x_j are co-associated. However, the semantics of a co-association are equivocal as the semantics of an association are ambiguous by itself (e.g., author vs. editor). Therefore, instead of relying solely on document associations, we use the textual data of associated documents to construct an entity representation.

Vector space models for document retrieval, such as LSI [70] or LDA [38], can be adapted to entity retrieval. We substantiate this for a specific entity retrieval task: expert finding. As there are many more documents than experts, it is not ideal to estimate a vector space directly on the expert-level using bag-of-word vectors (e.g., by representing every expert as a concatenation of its documents) due to data sparsity. Therefore, it is preferable to first estimate a vector space on the document collection and then use the obtained document representations to construct an entity vector. Demartini et al. [71] take an entity's representation to be the sum of its documents:

$$e_i = \sum_{d_j \in D_{x_i}} g(d_j), \quad (6.1)$$

where e_i is the k -dimensional vector representation of entity $x_i \in X$ and g is the function mapping a document to its vector space representation (e.g., LSI). The dimensionality k depends on the underlying vector space. For simple bag-of-words representations, k is equal to the number of words in the vocabulary. For latent vector spaces (e.g., LSI), the k -dimensional space encodes latent concepts and the choice of k is left to the user.

Vector space models for document retrieval are often constructed heuristically. E.g., Eq. 6.1 does not make optimal use of document-entity associations as document representations are added without taking into consideration the significance of words contained within them [154]. And if many diverse documents are associated with an expert, then Eq. 6.1 is likely to succumb to the noise in these vectors and yield meaningless representations.

To address this problem, Le and Mikolov [141] introduced doc2vec by adapting the word2vec models to incorporate the document a phrase occurs in. They optimize word

and document embeddings jointly to predict a word given its context and the document the word occurs in. The key difference between `word2vec` and `doc2vec` is that the latter considers an additional meta-token in the context that represents the document. Instead of performing dimensionality reduction on bag-of-words representations, `doc2vec` learns representations from word phrases. Therefore, we use the `doc2vec` model to learn expert embeddings by representing every expert $x_j \in X$ as a pseudo-document consisting of the concatenation of their associated documents D_{x_j} .

A different neural language model architecture than `doc2vec` was proposed by Van Gysel et al. [257], specifically for the expert finding task. For a given word w_i and expert x_j :

$$\text{score}(w_i, x_j) = e^{(v_i^T \cdot e_j + b_j)}, \quad (6.2)$$

where v_i (e_j , resp.) are the latent k -dimensional representations of word w_i (and expert x_j , respectively) and b_j is the bias scalar associated with expert x_j . Eq. 6.2 can be interpreted as the unnormalized factor product of likelihood $P(w_i | x_j)$ and prior $P(x_j)$ in log-space. The score is then transformed to the conditional probability

$$P(X = x_j | w_i) = \frac{\text{score}(w_i, x_j)}{\sum_{x_l \in X} \text{score}(w_i, x_l)}.$$

Unlike Eq. 6.1, the conditional probability distribution $P(X = x_j | w_i)$ will be skewed towards relevant experts if the word w_i is significant as described by Luhn [154]. The parameters v_i , e_j and b_j are learned from the corpus using gradient descent. See [257] for details.

Our focus lies on representations of entities e_j and how these correspond to structures within their domains (i.e., organizations for experts). These representations are estimated using a corpus only and can be interpreted as vectors in word embedding space that correspond to entities (i.e., people) instead of words.

6.4 Experimental set-up

6.4.1 Research questions

We investigate regularities within text-based entity vector spaces, using expert finding as our concrete test case, and ask how these representations correspond to structure in their respective domains. As indicated in the introduction of this chapter, we seek to answer the following research questions:

RQ2.5 Do clusterings of text-based entity representations reflect the structure of their domains?

Many organizations consist of smaller groups, committees or teams of experts who are appointed with a specific role. When we cluster expert representations, do the clusters correspond to these groups?

RQ2.6 To what extent do different text-based entity representation methods encode relations between entities?

The associations within expert domains encode a co-association graph structure. To what extent do the different expertise models encode this co-association between experts? In particular, if we rank experts according to their nearest neighbours, how does this ranking correspond to the academic co-author graph?

6.4.2 Expert finding collections

We use a subset of the expert finding collections of Chapter 5: the publicly-available expert finding collections provided by the World Wide Web Consortium (W3C) and Tilburg University (TU). We refer to Table 5.1 (previous chapter) for an overview. Note that we do not use the CSIRO Enterprise Research Collection (CERC) in this chapter due to the lack of information about the structure of the organization.

W3C The W3C collection was released as part of the 2005–2006 editions of the TREC Enterprise Track [64]. It contains a heterogeneous crawl of W3C’s website (June 2004) and consists of mailing lists and discussion boards among others. In the 2005 edition, TREC released a list of working groups and their members. Each working group is appointed to study and report on a particular aspect of the World Wide Web to enable the W3C to pursue its mission. We use the associations provided by Van Gysel et al. [257], which they gathered by applying named entity recognition and linking these mentions to a list of organization members, as proposed by Balog et al. [20].

TU The TU collection consists of a crawl of a university’s internal website and contains bi-lingual documents, such as academic publications, course descriptions and personal websites [36]. The document-candidate associations are part of the collection. For every member of the academic staff, their academic title is included as part of the collection.

6.4.3 Implementations and parameters

We follow a similar experimental set-up as previous work [20, 71, 176, 257]. For LSI, LDA, word2vec and doc2vec we use the Gensim¹ implementation, while for the log-linear model we use the Semantic Entity Retrieval Toolkit² (SERT) [261] that was released as part of Chapter 5 and is more closely described in Appendix B.

The corpora are normalized by lowercasing and removing punctuation and numbers. The vocabulary is pruned by removing stop words and retaining the 60k most frequent words. We sweep exponentially over the vector space dimensionality ($k = 32, 64, 128$ and 256) of the methods under comparison. This allows us to evaluate the effect of differently-sized vector spaces and their modelling capabilities.

For word2vec, a query/document is represented by its average word vector, which is effective for computing short text similarity [128]. We report both on the Continuous Bag-of-Words (CBOW) and Skip-gram (SG) variants of word2vec.

¹<https://radimrehurek.com/gensim>

²<https://github.com/cvangysel/SERT>

For LDA, we set $\alpha = \beta = 0.1$ and train the model for 100 iterations or until topic convergence is achieved. Default parameters are used in all other cases. Unlike Van Gysel et al. [257], we do not initialize with pre-trained word2vec embeddings.

For LSI, LDA and word2vec, expert representations are created from document representations according to Eq. 6.1.

In addition to text-based representations, we also include two baselines that do not consider textual data. For the first method (Graph PCA), we construct a weighted, undirected co-association graph where the weight between two entities is given by the number of times they are co-associated. We then apply Principal Component Analysis to create a latent representation for every entity. Secondly, we include a baseline where experts are represented as a random vector sampled from a standard multivariate normal distribution.

6.5 Regularities in entity vector spaces

We investigate regularities within latent text-based entity vector spaces. In particular, we first build latent representations for experts and ground these in the structure of the organizations where these experts are active. First, we cluster latent expert representations using different clustering techniques and compare the resulting clusters to committees in a standards organization of the World Wide Web (**RQ2.5**). We continue by investigating to what extent these representations encode entity relations (**RQ2.6**). We complement the answers to our research questions with an analysis of the prior (the scalar bias in Eq. 6.2) associated with every expert in one of the models we consider, SERT, and compare this to their academic rank.

6.5.1 Answers to research questions

RQ2.5 Do clusterings of text-based entity representations reflect the structure of their domains?

The World Wide Web Consortium (W3C) consists of various working groups.³ Each working group is responsible for a particular aspect of the WWW and consists of two or more experts. We use these working groups as ground truth for evaluating the ability of expert representations to encode similarity. The W3C working groups are special committees that are established to produce a particular deliverable [214, p. 492] and are a way to gather experts from around the organization who share areas of expertise and who would otherwise not directly communicate. Working groups are non-hierarchical in nature and represent clusters of experts. Therefore, they can be used to evaluate to what extent entity representations can be used as feature vectors for clustering.

We cluster expert representations using K -means [162]. While K -means imposes strong assumptions on cluster shapes (convexity and isotropism), it is still very popular today due to its linear time complexity, geometric interpretation and absence of hard to choose hyper-parameters (unlike spectral variants or DBSCAN). We cluster expert

³<http://www.w3.org/Consortium/activities>

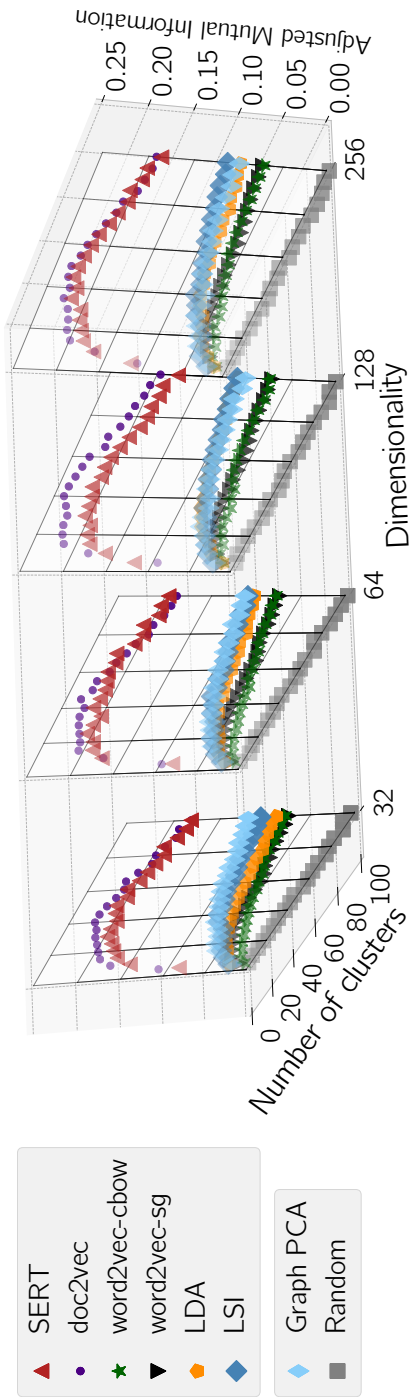


Figure 6.1: Comparison of clustering capabilities of expert representations (random, Graph PCA, LSI, LDA, word2vec, doc2vec and SERT) using K -means for $10^0 \leq K \leq 10^2$ (y-axis). The x-axis shows the dimensionality of the representations and the z-axis denotes the Adjusted Mutual Information.

representations of increasing dimensionality k ($k = 2^i$ for $5 \leq i < 9$) using a linear sweep over the number of clusters K ($10^0 \leq K < 10^2$).

During evaluation we transform working group memberships to a hard clustering of experts by assigning every expert to the smallest working group to which they belong as we wish to find specialized clusters contrary to general clusters that contain many experts. We then use Adjusted Mutual Information, an adjusted-for-chance variant of Normalized Information Distance [268], to compare both clusterings. Adjusting for chance is important as non-adjusted measures (such as BCubed precision/recall⁴ as presented by Amigó et al. [9]) have the tendency to take on a higher value for a larger value of K . Performing the adjustment allows us to compare clusterings for different values of K . We repeat the K -means clustering 10 times with different centroids initializations and report the average.

Fig. 6.1 shows the clustering capabilities of the different representations for different values of K and vector space dimensionality. Ignoring the random baseline, representations built using word2vec perform worst. This is most likely due to the fact that document representations for word2vec are constructed by averaging individual word vectors. Next up, we observe a tie between LSI and LDA. Interestingly enough, the baseline that only considers entity-document associations and does not take into account textual content, Graph PCA, outperforms all representations constructed from document-level vector space models (Eq. 6.1). Furthermore, doc2vec and SERT perform best, regardless of vector space dimensionality, and consistently outperform the other representations. If we look at the vector space dimensionality, we see that the best clustering is created using 128-dimensional vector spaces. Considering the number of clusters, we see that doc2vec and SERT peak at about 40 to 60 clusters. This corresponds closely to the number of ground-truth clusters. The remaining representations (word2vec, LSI, LDA, Graph PCA) only seem to plateau in terms of clustering performance at $K = 100$, far below the clustering performance of the doc2vec and SERT representation methods.

To answer our first research question, we conclude that expert representations can be used to discover structure within organizations. However, the quality of the clustering varies greatly and use of more advanced methods (i.e., doc2vec or SERT) is recommended.

RQ2.6 To what extent do different text-based entity representation methods encode relations between entities?

The text-based entity representation problem is characterized by a bipartite graph of entities and documents where an edge denotes an entity-document association. This differs from entity finding settings where explicit entity-entity relations are available and fits into the scenario where representations have to be constructed from unstructured text only. If latent text-based entity representations encode co-associations, then we can use this insight for (1) a better understanding of text-based entity representation models, and (2) the usability of latent text-based entity representations as feature vectors in scenarios where relations between entities are important.

⁴This can be verified empirically by computing BCubed measures for an increasing number of random partitions.

Table 6.1: Retrieval performance (NDCG and R-Precision) when ranking experts for a query expert by the cosine similarity of expert representations (random, Graph PCA, LSI, LDA, word2vec, doc2vec and SERT) for the TU expert collection (Section 6.4.2) for an increasing representation dimensionality. The relevance labels are given by the number of times two experts were co-authors of academic papers. Significance of results is determined using a two-tailed paired Student t-test (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$) between the best performing model and second best performing method.

Dimensionality $k =$	32		64		128		256	
	NDCG	R-Precision	NDCG	R-Precision	NDCG	R-Precision	NDCG	R-Precision
Random	0.18	0.01	0.18	0.01	0.18	0.01	0.18	0.01
Graph PCA	0.38	0.18	0.39	0.20	0.41	0.23	0.39	0.23
LSI	0.39	0.17	0.43	0.21	0.46	0.23	0.47	0.23
LDA	0.44	0.19	0.45	0.20	0.46	0.22	0.52	0.28
word2vec-sg	0.46	0.22	0.49	0.24	0.49	0.24	0.50	0.25
word2vec-cbow	0.46	0.23	0.47	0.24	0.48	0.25	0.48	0.25
doc2vec	0.35	0.14	0.36	0.15	0.36	0.16	0.35	0.15
SERT	0.53***	0.29***	0.54***	0.31***	0.53***	0.30***	0.53	0.31*

We evaluate the capacity of text-based expert representations to encode co-associations by casting the problem as a ranking task. Contrary to typical expert finding, where we rank experts according to their relevance to a textual query, for the purpose of answering **RQ2.6**, we rank experts according to their cosine similarity w.r.t. a query expert [18]. This task shares similarity with content-based recommendation based on unstructured data [198].

In expert finding collections, document-expert associations can indicate many things. For example, in the W3C collection, entity-document associations are mined from expert mentions [20]. However, for the TU collection, we know that a subset of associations corresponds to academic paper authorship. Therefore, we construct ranking ground-truth from paper co-authorship and take the relevance label of an expert to be the number of times the expert was a co-author with the query expert (excluding the query expert themselves). Our intention is to determine to what extent latent entity representations estimated from text can reconstruct the original co-author graph. Given that we estimate the latent entity representations using the complete TU document collection, by design, our evaluation is contained within our training set for the purpose of this analysis.

Table 6.1 shows NDCG and R-Precision [165, p. 158] for various representation models and dimensionality. SERT performs significantly better than the other representations methods (except for the 256-dimensional representations where significance was not achieved w.r.t. LDA). SERT is closely followed by word2vec (of which both variants score only slightly worse than SERT), LDA and LSI. The count-based distributional methods (LSI, LDA) perform better as the dimensionality of the representations increases. This is contrary to SERT, where retrieval performance is very stable across dimensionalities. Interestingly, doc2vec performs very poorly at reconstructing the co-author graph and is even surpassed by the Graph PCA baseline. This is likely due to the fact that doc2vec is trained on expert profiles and is not explicitly presented with document-expert associations. The difference in performance between doc2vec and SERT for **RQ2.6** reflects a difference in architecture: while SERT is directly optimized to discriminate between entities, doc2vec models entities as context in addition to language. Hence, similarities and dissimilarities between entities are preserved much better by SERT.

We answer our second research question as follows. Latent text-based entity representations do encode information about entity relations. However, there is a large difference in the performance of different methods. SERT seems to encode the entity co-associations better than other methods, by achieving the highest performance independent of the vector space dimensionality.

6.5.2 Analysis of the expert prior in the log-linear model

One of the semantic models that we consider, SERT, learns a prior $P(X)$ over entities. The remaining representation learning methods do not encode an explicit entity prior. It might be possible to extract a prior from generic entity vector spaces, e.g., by examining the deviation from the mean representation for every entity. However, developing such prior extraction methods are a topic of study by themselves and are out of scope for this chapter.

In the case of expert finding, this prior probability encodes a belief over experts




Figure 6.2: Box plots of prior probabilities learned by SERT, grouped by the experts' academic rank, for the TU collection. We only show the prior learned for a SERT model with $k = 32$, as the distributions of models with a different representation dimensionality are qualitatively similar.

without observing any evidence (i.e., query terms in SERT). Which structural information does this prior capture? We now investigate the regularities encoded within this prior and link it back to the hierarchy among scholars in the Tilburg University collection. We estimate a SERT model on the whole TU collection and extract the prior probabilities:

$$P(X = x_i) = \frac{e^{(b_i)}}{\sum_l e^{(b_l)}}, \quad (6.3)$$

where b is the bias vector of the SERT model in Eq. 6.2.

For 666 out of 977 experts in the TU collection we have ground truth information regarding their academic rank [36].⁵ Fig. 6.2 shows box plots of the prior probabilities, learned automatically by the SERT model from only text and associations, grouped by academic rank. Interestingly, the prior seems to encode the hierarchy amongst scholars at Tilburg University, e.g., Post-docs are ranked higher than PhD students. This is not surprising as it is quite likely that higher-ranked scholars have more associated documents.

The prior over experts in SERT encodes rank within organizations. This is not surprising, as experts (i.e., academics in this experiment) of higher rank tend to occur more frequently in the expert collection. This observation unveils interesting insights about the expert finding task and consequently models targeted at solving it. Unlike unsupervised ad-hoc document retrieval where we assume a uniform prior and normalized document lengths, the prior over experts in the expert finding task is of much greater importance. In addition, we can use this insight to gain a better understanding of the formal language models for expertise retrieval [20]. Balog et al. [20] find that, for the expert finding task, the document-oriented language model performs better than an entity-oriented language model. However, the document-oriented model [20] will rank experts with more associated documents higher than experts with few associated documents. On the contrary, the entity-oriented model of Balog et al. [20], imposes a uniform prior over experts. SERT is an entity-oriented model and performs better than the formal document-oriented language model [257]. This is likely due to the fact that SERT learns an empirical prior over entities instead of making an assumption of uniformity, in addition to its entity-oriented perspective.

In the case of general entity finding, the importance of the number of associated documents might be of lesser importance. Other sources of prior information, such as link analysis [197], recency [95] and user interactions [226], can be a better way of modelling entity importance than the length of entity descriptions.

6.6 Summary

In this chapter we have investigated the structural regularities contained within latent text-based entity representations. Entity representations were constructed from expert finding collections using methods from distributional semantics (LSI), topic models (LDA) and neural networks (word2vec, doc2vec and SERT). For LSI, LDA

⁵ 126 PhD Students, 49 Postdoctoral Researchers, 210 Assistant Professors, 89 Associate Professors and 190 Full Professors; we filtered out academic ranks that only occur once in the ground-truth, namely Scientific Programmer and Research Coordinator.

and word2vec, document-level representations were transformed to the entity scope according to the framework of Demartini et al. [71]. In the case of doc2vec and SERT, entity representations were learned directly. In addition to representations estimated only from text, we considered non-textual baselines, such as: (1) random representations sampled from a Normal distribution, and (2) the rows of the dimensionality-reduced adjacency matrix of the co-association graph.

We have found that text-based entity representations can be used to discover groups inherent to an organization. We have clustered entity representations using K -means and compared the obtained clusters with a ground-truth partitioning. No information about the organization is presented to the algorithms. Instead, these regularities are extracted by the documents associated with entities and published within the organization. Furthermore, we have evaluated the capacity of text-based expert representations to encode co-associations by casting the problem as a ranking task. We discover that text-based representations retain co-associations up to different extents. In particular, we find that SERT entity representations encode the co-association graph better than the other representation learning methods. We conclude that this is due to the fact that SERT representations are directly optimized to discriminate between entities. Lastly, we have shown that the prior probabilities learned by semantic models encode further structural information. That is, we find that the prior probability over experts (i.e., members of an academic institution), learned as part of a SERT model, encodes academic rank. In addition, we discuss the similarities between SERT and the document-oriented language model [20] and find that the document association prior plays an important role in expert finding.

Our findings have shown insight into how different text-based entity representation methods behave in various applications. In particular, we find that the manner in which entity-document associations are encoded plays an important role. That is, representation learning methods that directly optimize the representation of the entity seem to perform best. When considering different neural representation learning models (doc2vec and SERT), we find that their difference in architecture allows them to encode different regularities. doc2vec models an entity as context in addition to language, whereas SERT learns to discriminate between entities given their language. Thus, doc2vec can more adequately model the topical nature of entities, while SERT more closely captures the similarities and dissimilarities between entities. In the case of expert finding, we find that the amount of textual data associated with an expert is a principal measure of expert importance.

While the focus of this chapter was the analysis of the structural regularities contained within latent entity representations, the focus of the following chapters will be once again the modelling of latent vector spaces. In Chapter 7, we move from expert finding to a different entity retrieval problem: product search. Compared to expert finding, the product search scenario is characterized by a larger amount of retrievable entities. In addition, the amount of textual content per entity is much less than in the expert finding case.

Learning Latent Vector Spaces for Product Search

7.1 Introduction

Retail through online channels has become an integral part of consumers' lives [172]. In addition to using these online platforms that generate hundreds of billions of dollars in revenue [87], consumers increasingly participate in multichannel shopping where they research items online before purchasing them in brick-and-mortar stores. Search engines are essential for consumers to be able to make sense of these large collections of products available online [124]. In the case of directed searching (in contrast to exploratory browsing), users formulate queries using characteristics of the product they are interested in (e.g., terms that describe the product's category) [217]. However, it is widely known that there exists a mismatch between queries and product representations where both use different terms to describe the same concepts [146]. Thus, there is an urgent need for better semantic matching methods.

Product search is a particular example of the more general entity finding task that is increasingly being studied. Other entity finding tasks considered recently include searching for people [23], books [88] and groups [148]. Products are retrievable entities where every product is associated with a description and one or more user reviews. Therefore, we use the terms “product” and “entity” interchangeably in this chapter. However, there are two important differences between product search and the entity finding task as defined by de Vries et al. [69]. First, in entity finding one retrieves entities of a particular type from large broad coverage multi-domain knowledge bases such as Wikipedia [19, 69]. In contrast, product search engines operate within a single domain which can greatly vary in size. Second, user queries in product search consist of free-form text [217], as opposed to the semi-structured queries with additional type or relational constraints being used in entity finding [22, 69].

In this chapter we tackle the problem of discriminating between products based on the language (i.e., descriptions and reviews) they are associated with. Existing methods that are aimed at discriminating between entities based on textual data learn word representations using a language modelling objective or heuristically construct entity representations [71, 257]. Our approach directly learns two things: a unidirectional mapping between words and entities, as well as distributed representations of both

words and entities. It does so in an unsupervised and automatic manner such that words that are strongly evidential for particular products are projected nearby those products. While engineering of representations is important in information retrieval [17, 39, 48, 71, 95, 289], unsupervised joint representation learning of words and entities has not received much attention. We fill this gap. Our focus on learning representations for an end-to-end task such as *product search* is in contrast to the large volume of recent literature on word representation learning [248] that has a strong focus on upstream components such as distributional semantics [177, 199], parsing [59, 248] and information extraction [59, 248]. In addition, our focus on *unsupervised* representation learning is in contrast to recent entity representation learning methods [39, 289] that heavily depend on precomputed entity relationships and cannot be applied in their absence.

In recent years, significant progress has been made concerning semantic representations of entities. We point out three key insights on which we build: (1) Distributed representations [115] learned by discriminative neural networks reduce the curse of dimensionality and improve generalization. Latent features encapsulated by the model are shared by different concepts and, consequently, knowledge about one concept influences knowledge about others. (2) Discriminative approaches outperform generative models if enough training data is available [24, 191] as discriminative models solve the classification problem directly instead of solving a more general problem first [266]. (3) The unsupervised neural retrieval model we introduced in Chapter 5 [257] does not scale as they model a distribution over all retrievable entities; the approach is infeasible during training if the collection of retrievable entities is large.

Building on these insights, we introduce Latent Semantic Entities (LSE), a method that learns separate representations of words and retrievable objects jointly for the case where mostly unstructured documents are associated with the objects (i.e., descriptions and user reviews for products) and without relying on predefined relationships between objects (e.g., knowledge graphs). LSE learns to discriminate between entities for a given word sequence by mapping the sequence into the entity representation space. Contrary to heuristically constructed entity representations [71], LSE learns the relationship between words and entities directly using gradient descent. Unlike the model we proposed in Chapter 5 [257], we avoid computing the full probability distribution over entities; we do so by using noise-contrastive estimation.

The research questions we ask in this chapter towards answering **RQ3** are as follows:

RQ3.1 How do the parameters of LSE influence its efficacy?

RQ3.2 How does LSE compare to latent vector models based on LDA, LSI and word2vec?

RQ3.3 How does LSE compare to a smoothed language model that applies lexical term matching?

RQ3.4 What is the benefit of incorporating LSE as a feature in a learning-to-rank setting?




Figure 7.1: Illustrative example of how entities are ranked in vector space models w.r.t. a projected query. Query q is projected into entity space E using mapping f (black arrow) and entities (black crosses) are ranked according to their similarity in decreasing order.

7.2 Related work

We refer to Section 2.2 of our background chapter (Chapter 2). On the topic of product search (Section 2.2.2.2), Duan and Zhai [78] study the problem of learning query intent representation for structured product entities. They emphasize that existing methods focus only on the query space and overlook critical information from the entity space and the connection in between. We agree that modelling the connection between query words and entities and propagating information from the entity representations back to words is essential. In contrast to their work, we consider the problem of learning representations for entities based on their associations with unstructured documents.

Learning the representations of entities is not new, and the topic is covered in Section 2.2.2.3. In contrast to existing methods for entity representation learning, we model representations of words and entities jointly in separate spaces, in addition to a mapping from word to entity representations, in an unsupervised manner.

In addition, latent semantic retrieval models (Section 2.2.1) are also relevant to this chapter. With Chapter 5 we saw the introduction of an LSM for entity retrieval, with an emphasis on expert finding; we noted that training the parameters of the log-linear model becomes infeasible when the number of entities increases. In this chapter we mitigate this problem by considering only a random sample of entities as negative examples during training. This allows us to efficiently estimate model parameters in large product retrieval collections, which is not possible using the approach we introduced in Chapter 5 due to its requirement to compute a normalization constant over all entities.

In this chapter, we tackle the task of learning latent continuous vector representations for e-commerce products for the purpose of product search. The focus of this chapter lies in the language modelling and representation learning challenge. We learn distributed representations [115] of words and entities and a mapping between the two. At retrieval time, we rank entities according to the similarity of their latent representations to the projected representation of a query. Our model LSE is compared against existing entity-oriented latent vector representations that have been created using LSI, LDA and word2vec. We provide an analysis of model parameters and give insight in the quality of the joint representation space.




Figure 7.2: Schematic representation of the Latent Semantic Entities model for a single word w . Word embeddings W_v (e_v -dim. for $|V|$ words), entity embeddings W_e (e_E -dim. for $|X|$ entities) and the mapping from words to entities (e_E -by- e_v matrix W , e_E -dim. vector b) are learned using gradient descent.

7.3 Latent vector spaces for entity retrieval

We first introduce a generalized formalism and notation for entity-oriented latent vector space models. After that, in Section 7.3.2, we introduce Latent Semantic Entities, a latent vector space model that jointly learns representations of words, entities and a mapping between the two directly, based on the idea that entities are characterized by the words they are associated with and vice versa. Product representations are constructed based on the n-grams the products are likely to generate based on their description and reviews, while word representations are based on the entities they are associated with and the context they appear in. We model the relation between word and product representations explicitly so that we can predict the product representation for a previously unseen word sequence.

7.3.1 Background

We focus on a product retrieval setting in which a user wants to retrieve the most relevant products on an e-commerce platform. As in typical information retrieval scenarios, the user encodes their information need as a query q and submits it to a search engine. Product search queries describe characteristics of the product the user is searching for, such as a set of terms that describe the product's category [217].

Below, X denotes the set of entities that we consider. For every $x_i \in X$ we assume to have a set of associated documents D_{x_i} . The exact relation between the entity and its documents depends on the problem setting. In this chapter, entities are products [78, 193] and documents associated with these products are descriptions and product reviews.

Latent vector space models rely on a function $f : V^+ \rightarrow E$ that maps a sequence of words (e.g., a query q during retrieval) from a vocabulary V to an e_E -dimensional continuous entity vector space $E \subset \mathbb{R}^{e_E}$. Every entity $x_i \in X$ has a corresponding vector representation $e_i \in E$. Let $S_c : E \times E \rightarrow \mathbb{R}$ denote the cosine similarity between vectors in E . For a given query q , entities x_i are ranked in decreasing order of the cosine similarity between e_i and the query projected into the space of entities, $f(q)$. Fig. 7.1 illustrates how entities are ranked according to a projected query. For LSI, f is defined as the multiplication of the term-frequency vector representation of q with the

rank-reduced term-concept matrix and the inverse of the rank-reduced singular value matrix [70]. In the case of LDA, f becomes the distribution over topics conditioned on q [38]. This distribution is computed as the sum of the topic distributions conditioned on the individual words of q . In this chapter, the embedding f is learned; see Section 7.3.3 below.

Traditional vector space models operate on documents instead of entities. Demartini et al. [71] extend document-oriented vector spaces to entities by representing an entity as a weighted sum of the representations of their associated documents:

$$e_i = \sum_{d_j \in D_{x_i}} r_{i,j} f(d_j) \quad (7.1)$$

where $f(d_j)$ is the vector representation of d_j and $r_{i,j}$ denotes the relationship weight between document d_j and entity x_i . In this chapter we put $r_{i,j} = 1$ whenever $d_j \in D_{x_i}$ for a particular $x_i \in X$ and $r_{i,j} = 0$ otherwise, as determining the relationship weight between entities and documents is a task in itself.

7.3.2 Latent semantic entities

While Eq. 7.1 adapts document-oriented vector space models to entities, in this chapter we define f by explicitly learning (Section 7.3.3) the mapping between word and entity representations and the representations themselves:

$$f(s) = \tanh \left(W \cdot (W_v \cdot \frac{1}{|s|} \sum_{w_i \in s} \delta_i) + b \right) \quad (7.2)$$

for a string s of constituent words $w_1, \dots, w_{|s|}$ (an n -gram extracted from a document or a user-issued query), where W_v is the $e_V \times |V|$ projection matrix that maps the averaged one-hot representations (i.e., a $|V|$ -dimensional vector with element i turned on and zero elsewhere) of word w_i , δ_i , to its e_V -dimensional distributed representation. This is equivalent to taking the embeddings of the words in s and averaging them. In addition, b is a e_E -dimensional bias vector, W is the $e_E \times e_V$ matrix that maps averaged word embeddings to their corresponding position in entity space E and \tanh is the element-wise smooth hyperbolic tangent with range $(-1, 1)$. This transformation allows word embeddings and entity embeddings to be of a different dimensionality.

In other words, for a given string of words we take the representation of this string to be the average of the representations of the words it contains [141, 177]. This averaged word representation is then transformed using a linear map (W) and afterwards translated using b . We then apply the hyperbolic tangent as non-linearity such that every component lies between -1 and 1 . First of all, this regularizes the domain of the space and avoids numerical instability issues that occur when the magnitude of the vector components becomes too large. Secondly, by making the function non-linear we are able to model non-linear class boundaries in the optimization objective that we introduce in the next section. We use W_e to denote the $|X| \times e_E$ matrix that holds the entity representations. Row i of W_e corresponds to the vector representation, e_i , of entity x_i . Fig. 7.2 depicts a schematic overview of the proposed model. The parameters

W_v , W , b and W_e will be learned automatically using function approximation methods as explained below.

The model proposed in this section shares similarities with previous work on word embeddings and unsupervised neural retrieval models [177, 257]. However, its novelty lies in its ability to scale to large collections of entities and its underlying assumption that words and entities are embedded in spaces of different dimensionality: (1) The model of [177] has no notion of entity retrieval as it estimates a language model for the whole corpus. (2) Similar to [177], Eq. 7.2 aggregates words $w_i \in s$ to create a single phrase representation of s . However, in [257], a distribution $P(X | w_i)$ is computed for every w_i independently and aggregation occurs using the factor product. This is infeasible during model training when the collection of retrievable objects becomes too large, as is the case for product search. In the next section (Section 7.3.3) we solve this problem by sampling. (3) In both [177, 257] two sets of representations of the same dimensionality are learned for different types of objects with potentially different latent structures (e.g., words, word contexts and experts). As mentioned earlier, Eq. 7.2 alleviates this problem by transforming one latent space to the other.

7.3.3 Parameter estimation

For a particular document $d \in D_{x_i}$ associated with entity x_i , we generate n-grams $w_{j,1}, \dots, w_{j,n}$ where n (window size) remains fixed during training. For every n-gram $w_{j,1}, \dots, w_{j,n}$, we compute its projected representation $f(w_{j,1}, \dots, w_{j,n})$ in E using f (Eq. 7.2). The objective, then, is to directly maximize the similarity between the vector representation of the entity e_i and the projected n-gram $f(w_{j,1}, \dots, w_{j,n})$ with respect to S_c (Section 7.3.1), while minimizing the similarity between $f(w_{j,1}, \dots, w_{j,n})$ and the representations of non-associated entities. This allows the model to learn relations between neighbouring words in addition to the associated entity and every word.

However, considering the full set of entities for the purpose of discriminative training can be costly when the number of entities $|X|$ is large. Therefore, we apply a variant of Noise-Contrastive Estimation (NCE) [106, 176, 183, 184] where we sample negative instances from a noise distribution with replacement. We use the uniform distribution over entities as noise distribution. Define

$$P(\mathcal{S} | e_i, f(w_{j,1}, \dots, w_{j,n})) = \sigma(e_i \cdot f(w_{j,1}, \dots, w_{j,n})) \quad (7.3)$$

as the similarity of two representations in latent entity space, where

$$\sigma(t) = \frac{1}{1 + e^{-t}}$$

denotes the sigmoid function and \mathcal{S} is an indicator binary random variable that says whether x_i is similar to $f(w_{j,1}, \dots, w_{j,n})$.

We then approximate the probability of an entity x_i given an n-gram by randomly

sampling z contrastive examples:

$$\begin{aligned}
 & \log \tilde{P}(x_i \mid w_{j,1}, \dots, w_{j,n}) \\
 &= \log P(\mathcal{S} \mid e_i, f(w_{j,1}, \dots, w_{j,n})) \\
 &+ \sum_{\substack{k=1, \\ x_k \sim U(X)}}^z \log(1 - P(\mathcal{S} \mid e_k, f(w_{j,1}, \dots, w_{j,n})))
 \end{aligned} \tag{7.4}$$

where $U(X)$ denotes the uniform distribution over entities X , the noise distribution used in NCE [106]. Eq. 7.4 avoids iterating over all entities during parameter estimation as we stochastically sample z entities uniformly as negative training examples.¹

During model construction we maximize the log-probability (7.4) using batched gradient descent. The loss function for a single batch of m instances $((w_{k,1}, \dots, w_{k,n}), x_k)$ consisting of n -grams sampled from documents D_{x_k} (see Section 7.4.2) and associated entity x_k is as follows:

$$\begin{aligned}
 & L(W_v, W_e, W, b) \\
 &= -\frac{1}{m} \sum_{k=1}^m \log \tilde{P}(x_k \mid w_{k,1}, \dots, w_{k,n}) \\
 &+ \frac{\lambda}{2m} \left(\sum_{i,j} W_{v_{i,j}}^2 + \sum_{i,j} W_{e_{i,j}}^2 + \sum_{i,j} W_{i,j}^2 \right),
 \end{aligned} \tag{7.5}$$

where λ is a weight regularization parameter. Instances are shuffled before batches are created. The update rule for a particular parameter θ (W_v , W_e , W or b) given a single batch of size m is:

$$\theta^{(t+1)} = \theta^{(t)} - \alpha^{(t)} \odot \frac{\partial L}{\partial \theta}(W_v^{(t)}, W_e^{(t)}, W^{(t)}, b^{(t)}), \tag{7.6}$$

where $\alpha^{(t)}$ and $\theta^{(t)}$ denote the per-parameter learning rate and parameter θ at time t , respectively. The learning rate α consists of the same number of elements as there are parameters; in the case of a global learning rate, all elements of α are equal to each other. The derivatives of the loss function (7.5) are given in the Appendix.

7.4 Experimental setup

7.4.1 Research questions

In this chapter we investigate the problem of constructing a latent vector model of words and entities by directly modelling the discriminative relation between entities and word context. As indicated in the introduction of this chapter, we seek to answer the following research questions:

¹We exploit the special nature of our evaluation scenario where we know the unique association between documents and entities. The setup can easily be adapted to the more general case where a document is associated with multiple entities by extracting the same word sequences from the document for every associated entity.

RQ3.1 How do the parameters of LSE influence its efficacy?

In Section 7.3 we introduced various hyper-parameters along with the definition of Latent Semantic Entities. We have the size of word representations e_V and the dimensionality of the entity representations e_E . During parameter estimation, the window size n influences the context width presented as evidence for a particular entity. What is the influence of these parameters on the effectiveness of LSE and can we identify relations among parameters?

RQ3.2 How does LSE compare to latent vector models based on LDA, LSI and word2vec?

Is there a single method that always performs best or does effectiveness differ per domain? Does an increase in the vector space dimensionality impact the effectiveness of these methods?

RQ3.3 How does LSE compare to a smoothed language model that applies lexical term matching?

How does LSE compare to language models on a per-topic basis? Are there particular topics that work especially well with either type of ranker?

RQ3.4 What is the benefit of incorporating LSE as a feature in a learning-to-rank setting?

What if we combine popularity-based, exact matching and latent vector space features in a linear learning-to-rank setting? Do we observe an increase in effectiveness if we combine these features?

7.4.2 Experimental design

To answer the research questions posed in Section 7.4.1, we evaluate LSE in an entity retrieval setting organized around Amazon products (see Section 7.4.3). We choose to experiment with samples of Amazon product data [169, 170] for the following reasons: (1) The collection contains heterogeneous types of evidential documents associated with every entity: descriptions as well as reviews. (2) Every department (e.g., *Home & Kitchen*) constitutes a separate, self-contained domain. (3) Within each department there is a hierarchical taxonomy that partitions the space of entities in a rich structure. We can use the labels associated with these partitions and the partitions themselves as ground truth during evaluation. (4) Every department consists of a large number of products categorized over a large number of categories. Importantly, this allows us to construct benchmarks with an increasing number of entities. (5) Every product has a variety of attributes that can be used as popularity-based features in a learning-to-rank setting.

To answer **RQ3.1** we investigate the relation between the dimensionality of the entity representations e_E and window size n . The latter, the window size n , controls the context width the model can learn from, while the former, the dimensionality of the entity representations e_E , influences the number of parameters and expressive power of the model. We sweep exponentially over n (2^i for $0 \leq i < 6$) and e_E (2^i for $6 \leq i < 11$). **RQ3.2** is answered by comparing LSE with latent vector space model baselines (Section 7.4.5) for an increasing entity space dimensionality e_E (2^i

Table 7.1: Overview of the *Home & Kitchen*, *Clothing*, *Shoes & Jewelry*, *Pet Supplies* and *Sports & Outdoors* product search benchmarks. T and V denote the test and validation sets, respectively. Arithmetic mean and standard deviation are reported wherever applicable.

	Home & Kitchen	Clothing, Shoes & Jewelry	Pet Supplies	Sports & Outdoors
Corpus (train)				
Number of documents	88,130	94,024	416,993	502,313
Document length	70.02 ± 73.82	58.41 ± 61.90	77.48 ± 78.44	72.52 ± 81.47
Number of entities				
Documents per entity	8,192	16,384	32,768	65,536
	10.76 ± 52.01	5.74 ± 18.60	12.73 ± 55.98	7.66 ± 30.38
Topics (test)				
Topics	657 (T)	750 (T)	385 (T)	1,879 (T)
	72 (V)	83 (V)	42 (V)	208 (V)
Terms per topic	5.11 ± 1.79	4.10 ± 1.86	3.73 ± 1.62	4.64 ± 1.68
Relevant entities per topic	10.92 ± 32.41 (T)	20.15 ± 57.78 (T)	75.96 ± 194.44 (T)	29.27 ± 61.71 (T)
	10.29 ± 15.66 (V)	12.13 ± 19.85 (V)	57.40 ± 88.91 (V)	38.25 ± 157.34 (V)

for $6 \leq i < 11$). For **RQ3.3**, we compare the per-topic paired differences between LSE and a lexical language model. In addition, we investigate the correlation between lexical matches in relevant entity documents and ranker preference. We address **RQ3.4** by evaluating LSE as a feature in a machine-learned ranking in addition to query-independent and lexical features.

The number of n -grams sampled per entity $x \in X$ from associated documents D_x in every epoch (i.e., iteration of the training data) is equal to

$$\left\lceil \frac{1}{|X|} \sum_{d \in D} \max(|d| - n + 1, 0) \right\rceil,$$

where the $|\cdot|$ operator is used interchangeably for the size of set X and the number of tokens in documents $d \in D$. This implicitly imposes a uniform prior over entities (i.e., stratified sampling where every entity is of equal importance). The word vocabulary V is created for each benchmark by ignoring punctuation, stop words and case; numbers are replaced by a numerical placeholder token. We prune V by only retaining the 2^{16} most-frequent words so that each word can be encoded by a 16-bit unsigned integer. In terms of parameter initialization of the Latent Semantic Entities model, we sample the initial matrices W_v , W (Eq. 7.2) and W_e uniformly in the range $\left[-\sqrt{\frac{6.0}{m+n}}, \sqrt{\frac{6.0}{m+n}}\right]$ for an $m \times n$ matrix, as this initialization scheme is known to improve model training convergence [91], and take the bias vector b to be null. The number of word features is set to $e_V = 300$, similar to [176]. We take the number of negative examples $z = 10$ to be fixed. Mikolov et al. [176] note that a value of z between 10 and 20 is sufficient for large data sets [183].

We used Adam ($\alpha = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$) [133] with batched gradient descent ($m = 4096$) and weight decay $\lambda = 0.01$ during training on NVidia Titan X GPUs. Adam has been designed specifically for non-stationary, stochastic cost functions like the one we defined in Eq. 7.4. For every model, we iterate over the training data 15 times and choose the best epoch based on the validation sets (Table 7.1).

7.4.3 Product search benchmarks

We evaluate on four samples from different product domains² (Amazon departments), each with of an increasing number of products: *Home & Kitchen* (8,192 products), *Clothing, Shoes & Jewelry* (16,384 products), *Pet Supplies* (32,768 products) and *Sports & Outdoors* (65,536 products); see Table 7.1. The documents associated with every product consist of the product description plus reviews provided by Amazon customers.

Rowley [217, p. 24] describes directed product search as users searching for “a producer’s name, a brand or a set of terms which describe the category of the product.” Following this observation, the test topics c_i are extracted from the categories each product belongs to. Category hierarchies of less than two levels are ignored, as the first level in the category hierarchy is often non-descriptive for the product (e.g., in *Clothing, Shoes & Jewelry* this is the gender for which the clothes are designated). Products belonging to a particular category hierarchy are considered as relevant for its extracted topic. Products can be relevant for multiple topics. Textual representations q_{c_i} of the topics based on the categories are extracted as follows. For a single hierarchy of categories, we tokenize the titles of its sub-categories and remove stopwords and duplicate words. For example, a digital camera lense found in the *Electronics* department under the categorical topic *Camera & Photo* \rightarrow *Digital Camera Lenses* will be relevant for the textual query “*photo camera lenses digital*.” Thus, we only have two levels of relevance. We do not index the categories of the products as otherwise the query would match the category and retrieval would be trivial.

7.4.4 Evaluation measures and significance

To measure retrieval effectiveness, we report Normalized Discounted Cumulative Gain (NDCG). For **RQ3.4**, we additionally report Precision@k ($k = 5, 10$). Unless mentioned otherwise, significance of observed differences is determined using a two-tailed paired Student’s t-test [234] (** $p < 0.01$; * $p < 0.05$; $p < 0.1$).

7.4.5 Methods used in comparisons

We compare Latent Semantic Entities to state-of-the-art latent vector space models for entity retrieval that are known to perform semantic matching [146]. We also conduct a contrastive analysis between LSE and smoothed language models with exact matching capabilities.

Vector Space Models for entity finding. Demartini et al. [71] propose a formal model for finding entities using document vector space models (Section 7.3.1). We compare the retrieval effectiveness of LSE with baseline latent vector space models created using (1) Latent Semantic Indexing (LSI) [70] with TF-IDF term weighting, (2) Latent Dirichlet Allocation (LDA) [38] with $\alpha = \beta = 0.1$, where a document is represented by its topic distribution, and (3) word2vec [177] with CBOW and negative sampling, where a query/document is represented by the average of its word embeddings

²A list of product identifiers, topics and relevance assessments can be found at <https://github.com/cvangysel/SERT>.

(same for queries in LSE). Similar to LSE, we train word2vec for 15 iterations and select the best-performing model using the validation sets (Table 7.1).

Query-likelihood Language Model. For every entity a profile-based statistical language model is constructed using maximum-likelihood estimation [20, 152, 266], which is then smoothed by the language model of the entire corpus. The retrieval score of entity x for query q is defined as

$$\tilde{P}(q | x) = \prod_{t_i \in q} P(t_i | \theta_x), \quad (7.7)$$

where $P(t | \theta_x)$ is the probability of term t occurring in the smoothed language model of x (Jelinek-Mercer smoothing [286]). Given a query q , entities are ranked according to $\tilde{P}(q | x)$ in descending order.

Machine-learned ranking. RankSVM models [125] in Section 7.5.2 and 7.6 are trained using stochastic gradient descent using the implementation of Sculley and Inc [227]. We use default values for all parameters, unless stated otherwise. For the experiment investigating LSE as a feature in machine-learned ranking in Section 7.5.2, we construct training examples by using the relevant entities as positive examples. Negative instances are generated by sampling from the non-relevant entities with replacement until the class distribution is uniform.

7.5 Results & discussion

We start by giving a high-level overview of our experimental results (**RQ3.1** and **RQ3.2**), followed by a comparison with lexical matching methods (**RQ3.3**) and the use of LSE as a ranking feature (**RQ3.4**) (see Section 7.4.2 for an overview of the experimental design).

7.5.1 Overview of experimental results

RQ3.1 Fig. 7.3 depicts a heat map for every combination of window size and entity space dimensionality evaluated on the validation sets (Table 7.1).

Fig. 7.3 shows that neither extreme values for the dimensionality of the entity representations nor the context width alone achieve the highest performance on the validation sets. Instead, a low-dimensional entity space (128- and 256-dimensional) combined with a medium-sized context window (4- and 8-grams) achieve the highest NDCG. In the two largest benchmarks (Fig. 7.3c, 7.3d) we see that for 16-grams, NDCG actually lowers as the dimensionality of the entity space increases. This is due to the model *fitting* the optimization objective (Eq. 7.5), which we use as an unsupervised surrogate of relevance, too well. That is, as the model is given more learning capacity (i.e., higher dimensional representations), it starts to learn more regularities of natural language which counteract retrieval performance.

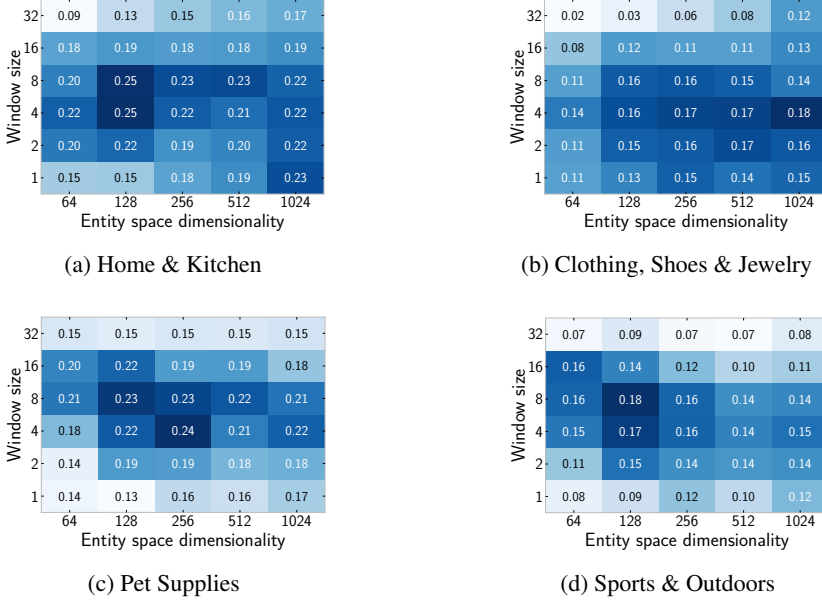


Figure 7.3: Sensitivity analysis of LSE in terms of NDCG for window size n and the size of entity representations e_E during parameter estimation (Eq. 7.5) for models trained on *Home & Kitchen*, *Clothing, Shoes & Jewelry*, *Pet Supplies* and *Sports & Outdoors* product search benchmarks (Section 7.4.3) and evaluated on the **validation** sets.

RQ3.2 Fig. 7.4 presents a comparison between LSE (window size $n = 4$) and vector space model baselines (Section 7.4.5) for increasing entity representation dimensionality (2^i for $6 \leq i < 11$) on the test sets.

LSE significantly outperforms ($p < 0.01$) all baseline methods in most cases (except for Fig. 7.4a where $e_E = 1024$). For the smaller benchmarks (Fig. 7.4a, 7.4b), we see LSI as the main competitor of LSE. However, as the training corpora become larger (in Fig. 7.4c, 7.4d), word2vec outperforms LSI and becomes the main contender of LSE. On all benchmarks, LSE peaks when the entity representations are low-dimensional (128- or 256-dimensional) and afterwards (for a higher dimensionality) performance decreases. On the other hand, word2vec stagnates in terms of NDCG around representations of 512 dimensions and never achieves the same level as LSE did for one or two orders of magnitude (base 2) smaller representations. This is a beneficial trait of LSE, as high-dimensional vector spaces are undesirable due to their high computational cost during retrieval [271].

7. Learning Latent Vector Spaces for Product Search

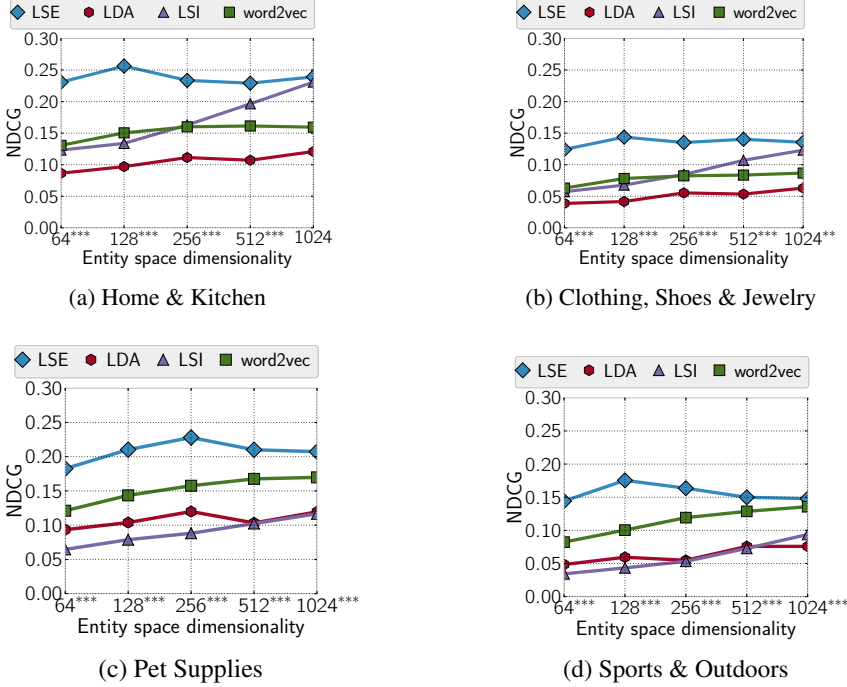


Figure 7.4: Comparison of LSE (with window size $n = 4$) with latent vector space baselines (LSI, LDA and word2vec; Section 7.4.5) on *Home & Kitchen*, *Clothing, Shoes & Jewelry*, *Pet Supplies* and *Sports & Outdoors* product search benchmarks (Section 7.4.3) and evaluated on the **test** sets. Significance (Section 7.4.4) is computed between LSE and the baselines for each vector space size.

7.5.2 A feature for machine-learned ranking

We now investigate the use of LSE as a feature in a learning to rank setting [151]. Latent vector space models are known to provide a means of semantic matching as opposed to a purely lexical matching [146, 257]. To determine to which degree this is indeed the case, we first perform a topic-wise comparison between LSE and a lexical language model, the Query-likelihood Language Model (QLM) [286], as described in Section 7.4.5. We optimize the parameters of LSE and QLM on the validation sets for every benchmark (Table 7.1). In the case of LSE, we select the model that performs best in Fig. 7.3. For QLM, we sweep over λ linearly from 0.0 to 1.0 (inclusive) with increments of 0.05.

RQ3.3 Fig. 7.5 shows the per-topic paired difference between LSE and QLM in terms of NDCG.

Topics that benefit more from LSE have a positive value on the y-axis, while those that prefer QLM have a negative value. We can see that both methods perform similarly for many topics (where $\Delta = 0.0$). For certain topics one method performs substantially better than the other, suggesting that the two are complementary. To further quantify

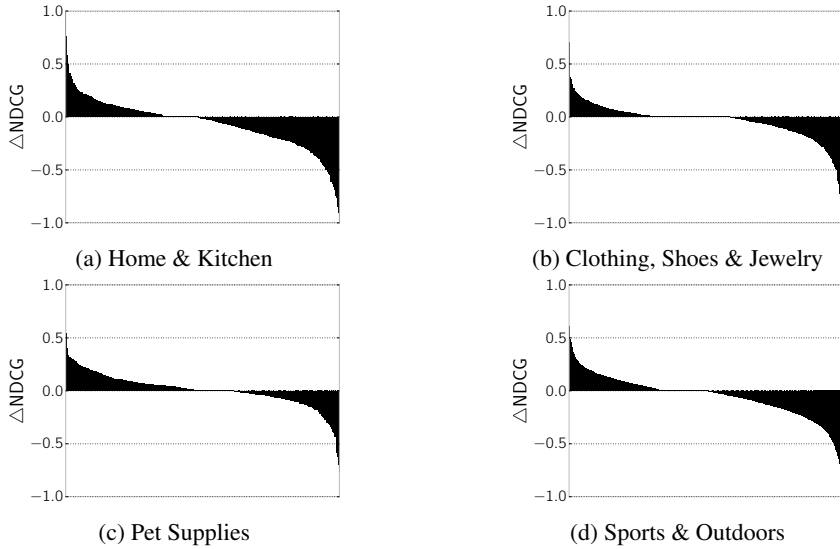


Figure 7.5: Per-topic paired differences between LSE and Query-likelihood Language Model for models trained on *Home & Kitchen*, *Clothing, Shoes & Jewelry*, *Pet Supplies* and *Sports & Outdoors* product search benchmarks (Section 7.4.3) and evaluated on the **test** sets. For every plot, the y-axis indicates ΔNDCG between LSE and a Query-likelihood Language Model. The x-axis lists the topics in the referenced benchmark in decreasing order of ΔNDCG such that topics for which LSE performs better are on the left and vice versa for the Query-likelihood Language Model on the right.

Table 7.2: Correlation coefficients between average IDF of lexically matched terms in documents associated with relevant entities and ΔNDCG . A negative correlation coefficient implies that queries consisting of more specific terms (i.e., low document freq.) that occur exactly in documents associated with relevant entities are more likely to benefit from QLM, whereas other queries (with less specific terms or less exact matches) gain more from LSE. Significance is achieved for all benchmarks ($p < 0.01$) using a permutation test.

Benchmark	Spearman R	Pearson R
Home & Kitchen	−0.30	−0.35
Clothing, Shoes & Jewelry	−0.40	−0.37
Pet Supplies	−0.17	−0.17
Sports & Outdoors	−0.34	−0.36

this, we investigate the relation between specific topic terms and their occurrence in documents relevant to these topics. That is, we measure the correlation between the per-topic ΔNDCG (as described above) and the average inverse document frequency (IDF) of exact/lexically matched terms in the profile-based language model. In Table 7.2

Table 7.3: Overview of the feature sets used in the machine-learned ranking experiments.

Features	Description
QI	Query-independent features:(1) product price; (2) product description length; (3) reciprocal of the Amazon sales rank; and (4) product PageRank scores based on four related product graphs (also bought, also viewed, bought together, buy after viewing).
QLM	Query-likelihood Language Model using Jelinek-Mercer smoothing with λ optimized on the validation set (Table 7.1). Posterior $P(q x)$ is used as a feature for entity x and query q .
LSE	Latent Semantic Entities optimized on the validation set (Table 7.1, Fig. 7.3). Similarity $S_c(f(q), e)$ is used as a feature for entity x , with vector representation e , and query q .

we observe that queries that contain specific tokens (i.e., with high inverse document frequency) and occur exactly in documents associated with relevant products, benefit more from QLM (lexical matches). Conversely, queries with less specific terms or without exact matches in the profiles of relevant products gain more from LSE (semantic matches).

This observation motivates the use of LSE as a ranking feature in addition to traditional language models. Specifically, we now evaluate the use of LSE as a feature in a linear RankSupport Vector Machine (SVM) (Section 7.4.5). Following Fang et al. [84], we consider query-independent (QI) popularity-based features in addition to features provided by LSE and QLM. This allows us to consider the effect of the query-dependent features independent from their ability to model a popularity prior over entities. Table 7.3 lists the feature sets.

RQ3.4 Table 7.4 shows the results for different combinations of feature sets used in a machine-learned ranker, RankSVM.

The experiment was performed using 10-fold cross validation on the test sets (Table 7.1). The combination using all features outperforms smaller subsets of features, on all metrics. We conclude that Latent Semantic Entities adds a signal that is complementary to traditional (lexical) language models, which makes it applicable in a wide range of entity-oriented search engines that use ranker fusion techniques.

7.6 Analysis of representations

Next, we analyse the entity representations e_i of the vector space models independent of the textual representations by providing empirical lower-bounds on their maximal retrieval performance, followed by a comparison with their actual performance so as to measure the effectiveness of word-to-entity mapping f .

Fig. 7.3 and 7.4 show which levels of performance may be achieved by using the latent models to generate a ranking from textual queries (Eq. 7.2). But this is only one

Table 7.4: Ranking performance results for query independent (QI) features, the Query-likelihood Language Model (QLM) match feature, the Latent Semantic Entities (LSE) match feature and combinations thereof, weighted using RankSVM (Section 7.5.2), evaluated on the test sets using 10-fold cross validation, for *Home & Kitchen*, *Clothing, Shoes & Jewelry*, *Pet Supplies* and *Sports & Outdoors* product search benchmarks (Section 7.4.3). The hyperparameters of the individual query features (QLM and LSE) were optimized using the validation sets. Significance of the results (Section 7.4.4) is computed between QI + QLM + LSE and QI + QLM.

	Home & Kitchen		
	NDCG	P@5	P@10
QI	0.005	0.002	0.001
QI + QLM	0.321	0.180	0.145
QI + LSE	0.257	0.121	0.107
QI + QLM + LSE	0.352***	0.192**	0.157***
	Clothing, Shoes & Jewelry		
	NDCG	P@5	P@10
QI	0.002	0.001	0.001
QI + QLM	0.177	0.079	0.068
QI + LSE	0.144	0.065	0.057
QI + QLM + LSE	0.198***	0.094***	0.080***
	Pet Supplies		
	NDCG	P@5	P@10
QI	0.003	0.002	0.002
QI + QLM	0.250	0.212	0.199
QI + LSE	0.268	0.222	0.214
QI + QLM + LSE	0.298***	0.255***	0.236***
	Sports & Outdoors		
	NDCG	P@5	P@10
QI	0.001	0.001	0.001
QI + QLM	0.235	0.183	0.156
QI + LSE	0.188	0.132	0.121
QI + QLM + LSE	0.264***	0.192***	0.172***

perspective. As entities are ranked according to their similarity with the projected query vector $f(q_c)$, the performance for retrieving entities w.r.t. the textual representation of a topic c depends on the structure of the entity space E , the ideal retrieval vector $e_c^* \in E$ (i.e., the vector that optimizes retrieval performance), and the similarity between $f(q_c)$ and e_c^* .

How can we determine the ideal vector e_c^* ? First, we define it to be the vector for which the cosine similarity with each of the entity embeddings results in a ranking where relevant entities are ranked higher than non-relevant or unjudged entities. We approximate e_c^* by optimizing the pair-wise SVM objective [125, 227]. That is, for every topic c we construct a separate RankSVM model based on its ground-truth as follows.




Figure 7.6: Comparison of the approximately ideal retrieval vector \tilde{e}_c^* with the projected query retrieval vector $f(q)$ for latent entity models built using LSE, LSI, LDA and word2vec (Section 7.4.5) on *Home & Kitchen* and *Pet Supplies* product search benchmarks (Section 7.4.3) and evaluated on the **test** sets. The plots for *Clothing*, *Shoes & Jewelry* and *Sports & Outdoors* product search benchmarks are qualitatively similar to the ones shown. The figures show the absolute performance in terms of NDCG of \tilde{e}_c^* (dashed curves) and $f(q)$ (solid curves); significance (Section 7.4.4) for the results for the approximately ideal retrieval vectors \tilde{e}_c^* is computed between LSE and the best-performing baseline for each vector space size and indicated along the x-axis.

We only consider topics with at least two relevant entities, as topics with a single relevant entity have a trivial optimal retrieval vector (the entity representation of the single relevant entity). Using the notation of [125], the normalized entity representations are used as features, and hence the feature mapping ϕ is defined as

$$\phi(c, x_i) = \frac{e_i}{\|e_i\|_2} \text{ for all } x_i \in X.$$

The target ranking r_c^* is given by the entities relevant to topic c . Thus, the features for every entity become the entity’s normalized representation and its label is positive if it is relevant for the topic and negative otherwise. The pair-wise objective then finds a weight vector such that the ranking generated by ordering according to the vector scalar product between the weight vector and the normalized entity representations correlates with the target ranking r_c^* . Thus, our approximation of the *ideal vector*, \tilde{e}_c^* , is given by the weight vector w_c for every c .³

What is the performance of this approximately ideal vector representation? And how far are our representations removed from it? Fig. 7.6 shows the absolute performance of \tilde{e}_c^* (dashed curves) and $f(q)$ (solid curves) in terms of NDCG. Comparing the (absolute) difference between every pair of dashed and solid curves for a single latent model gives an intuition of how much performance in terms of NDCG there is to gain by improving the projection function f for that method. The approximately ideal vectors \tilde{e}_c^* discovered for LSE outperform all baselines significantly. Interestingly, for representations created

³Note that \tilde{e}_c^* does not take into account the textual representations q_c of topic c , but only the clustering of entities relevant to c and their relation to other entities.

using LDA, the optimal performance goes up while the actual performance stagnates. This indicates that a higher vector space dimensionality renders better representations using LDA, however, the projection function f is unable to keep up in the sense that projected query vectors are not similar to the representations of their relevant entities. The latent models with the best representations (LSE and LSI) also have the biggest gap between $f(q)$ and \tilde{e}_c^* in terms of achieved NDCG.

We interpret the outcomes of our analysis as follows. The entity space E has more degrees of freedom to cluster entities more appropriately as the dimensionality of E increases. Consequently, the query projection function f is expected to learn a more complex function. In addition, as the dimensionality of E increases, so does the modelling capacity of the projection function f in the case of LSE and LSI (i.e., the transformation matrices become larger) and therefore more parameters have to be learned. We conclude that our method can more effectively represent entities in a lower-dimensional space than LSI by making better use of the vector space capacity. This is highly desirable, as the asymptotic runtime complexity of many algorithms operating on vector spaces increases at least linearly [271] with the size of the vectors.

7.7 Summary

We have introduced Latent Semantic Entities, an unsupervised latent vector space model for product search. It jointly learns a unidirectional mapping between, and latent vector representations of, words and products. We have also defined a formalism for latent vector space models where latent models are decomposed into a mapping from word sequences to the product vector space, representations of products in that space, and a similarity function. We have evaluated our model using Amazon product data, and compared it to state-of-the-art latent vector space models for product ranking (LSI, LDA and word2vec). LSE outperforms all baselines for lower-dimensional vector spaces.

In an analysis of the vector space models, we have compared the performance achieved with the ideal performance of the proposed product representations. We have shown that LSE constructs better product representations than any of the baselines. In addition, we have obtained important insights w.r.t. how much performance there is to gain by improving the individual components of latent vector space models. Future work can focus on improving the mapping from words to products by incorporating specialized features or increasing the mapping's complexity. In addition, semi-supervised learning may help specialize the vector space and mapping function for particular retrieval settings.

A comparison of LSE with a smoothed lexical language model unveils that the two methods make very different errors. Some directed product search queries require lexical matching, others benefit from the semantic matching capabilities of latent models. We have evaluated LSE as a feature in a machine-learned ranking setting and found that adding LSE to language models and popularity-based features significantly improves retrieval performance.

8

Neural Vector Spaces for Unsupervised Information Retrieval

8.1 Introduction

The vocabulary mismatch between query and document poses a critical challenge in search [146]. The vocabulary gap occurs when documents and queries, represented as a bag-of-words, use different terms to describe the same concepts. While improved semantic matching methods are urgently needed, in order for these methods to be effective they need to be applicable at early stages of the retrieval pipeline. Otherwise, candidate documents most affected by the mismatch (i.e., relevant documents that do not contain any query terms) will simply remain undiscovered. Boytsov et al. [43] show that (approximate) nearest neighbour algorithms [90, 188] can be more efficient than classical term-based retrieval. This strongly motivates the design of semantic matching methods that represent queries and documents in finite-dimensional vector spaces.

Latent semantic models, such as LSI [70], fit in the finite-dimensional vector space paradigm needed for nearest neighbour retrieval. However, LSI is known to retrieve non-relevant documents due to a lack of specificity [81]. The recent move towards learning word representations as part of neural language models [32] has shown impressive improvements in natural language processing (NLP) [59, 96, 177]. Therefore, it is reasonable to explore these representation learning methods for information retrieval (IR) as well. Unfortunately, in the case of full text document retrieval, only few positive results have been obtained so far [65]. We identify two causes for this shortfall.

First of all, IR tasks (e.g., document ranking) are fundamentally different from NLP tasks [65]. NLP deals with natural language regularities (e.g., discovering long range dependencies), whereas IR involves satisfying a user's information need (e.g., matching a query to a document). Therefore, specialized solutions and architectures for IR are needed.

Secondly, in the bag-of-words paradigm, query/document matching is performed by counting term occurrences within queries and documents. Afterwards, the frequencies are adjusted using weighting schemes that favor term specificity [215, 216], used as parameter values in probabilistic frameworks [286] and/or reduced in dimensionality [38, 70]. However, Baroni et al. [24] noted that for NLP tasks, prediction-based models, learned from scratch using gradient descent, outperform count-based models. Similarly,

the advances made by deep learning in computer vision [137] were not due to counting visual words [267] or by constructing complicated pipelines, but instead by optimizing a cost function using gradient descent and learning the model from scratch. However, for IR settings such as unsupervised news article retrieval, where one ranks documents in the absence of explicit or implicit relevance labels, prediction-based models have not received much attention. In order for deep learning to become feasible for unsupervised retrieval, we first need to construct an appropriate optimization objective.

In this chapter we introduce an optimization objective for learning latent representations of words and documents from scratch, in an unsupervised manner without relevance signals. Specifically, we introduce the Neural Vector Space Model (NVSM) for document retrieval. The optimization objective of NVSM mandates that word sequences extracted from a document should be predictive of that document. Learning a model of content in this way incorporates the following IR regularities: semantic matching (words occurring in each other's vicinity are learned to have similar representations), the clustering hypothesis (documents that contain similar language will have nearby representations in latent space), and term specificity (words associated with many documents will be neglected, as they have low predictive power).

One limitation of latent document vector spaces, including the NVSM we introduce here, is that their asymptotic complexity is bounded by the number of documents (i.e., one vector for every document) [4, 56, 256]. Consequently, the latent methods we consider in this chapter are only feasible to be constructed on document collections of medium scale. Therefore, we choose to evaluate our methods on article retrieval benchmarks ($\sim 200k$ to $500k$ documents each) from TREC [107, 286].

The research questions we ask in this chapter towards answering **RQ4** are as follows:

- RQ4.1** How does NVSM compare to other latent vector space models, such as doc2vec [141], word2vec [177, 270], LSI [70], LDA [38] and LSE [256], on the document retrieval task?
- RQ4.2** For what proportion of queries does NVSM perform better than the other rankers?
- RQ4.3** What gains does NVSM bring when combined with a lexical QLM and a competing state-of-the-art vector space model?
- RQ4.4** Do NVSMs exhibit regularities that we can link back to well-understood document collection statistics used in traditional retrieval models?

8.2 Related work

We refer to Section 2.2 of our background chapter (Chapter 2). Relevant to this chapter is the prior work on representation learning for document retrieval (Section 2.2.1) in addition to neural language modelling (Section 2.2.3) for natural language processing and automatic speech recognition.

The contribution of this chapter over and above the related work discussed earlier is the following. First, we substantially extend the LSE model introduced in the previous chapter (Chapter 7) and evaluate it on document search collections from TREC [244].

Our improvements over the LSE model are due to (1) increased regularization, and (2) accelerated training by reducing the internal covariate shift.

Second, we avoid *irrational exuberance* known to plague the deep learning field [65] by steering away from non-essential depth. That is, while we extend algorithms and techniques from deep learning, the model we present in this chapter is shallow.

Third, and contrary to previous work where information from pre-trained word embeddings is used to enrich the query/document representation of count-based models, our model, NVSM, is learned directly from the document collection without explicit feature engineering. We show that NVSM learns regularities known to be important for document retrieval from scratch. Our semantic vector space outperforms lexical retrieval models on some benchmarks. However, given that lexical and semantic models perform different types of matching, our approach is most useful as a supplementary signal to these lexical models.

8.3 Learning semantic spaces

In this section we provide the details of NVSM. First, we give a birds-eye overview of the model and its parameters and explain how to rank documents for a given query. Secondly, we outline our training procedure and optimization objective. We explain the aspects of the objective that make it work in a retrieval setting. Finally, we go into the technical challenges that arise when implementing the model and how we solved them in our open-source release.

8.3.1 The Neural Vector Space Model

Our work revolves around unsupervised ad-hoc document retrieval where a user wishes to retrieve documents (e.g., articles) as to satisfy an information need encoded in query q .

Below, a query q consists of terms (i.e., words) $t_1, \dots, t_{|q|}$ originating from a vocabulary V , where $|\cdot|$ denotes the length operator; D denotes the set of documents $\{d_1, \dots, d_{|D|}\}$. Every document $d_i \in D$ consists of a sequence of words $w_{i,1}, \dots, w_{i,|d_i|}$ with $w_{i,j} \in V$, for all $1 \leq j \leq |d_i|$.

We continue the work we started in Chapter 7 [256] and learn low-dimensional representations of words, documents and the transformation between them from scratch. That is, instead of counting term frequencies for use in probabilistic frameworks [286] or applying dimensionality reduction to term frequency vectors [70, 273], we learn representations directly by gradient descent from sampled n-gram/document pairs extracted from the corpus. These representations are embedded parameters in matrices $R_D \in \mathbb{R}^{|D| \times k_d}$ and $R_V \in \mathbb{R}^{|V| \times k_w}$ for documents D and vocabulary words V , respectively, such that $\vec{R}_V^{(i)}$ ($\vec{R}_D^{(j)}$, respectively) denotes the k_w -dimensional (k_d -dim., respectively) vector representation of word V_i (document d_i , respectively).

As the word representations $\vec{R}_V^{(i)}$ and document representations $\vec{R}_D^{(j)}$ are of different dimensionality, we require a transformation $f : \mathbb{R}^{k_w} \rightarrow \mathbb{R}^{k_d}$ from the word feature space to the document feature space. In this chapter we take the transformation to be

linear:

$$f(\vec{x}) = W\vec{x}, \quad (8.1)$$

where \vec{x} is a k_w -dimensional vector and W is a $k_d \times k_w$ parameter matrix that is learned using gradient descent (in addition to representation matrices R_V and R_D).

We compose a representation of a sequence of n words (i.e., an n -gram) w_1, \dots, w_n by averaging its constituent word representations:

$$g(w_1, \dots, w_n) = \frac{1}{n} \sum_{i=1}^n \vec{R}_V^{(w_i)}. \quad (8.2)$$

A query q is projected to the document feature space by the composition of f and g : $(f \circ g)(q) = h(q)$.

The matching score between a document d and query q is then given by the cosine similarity between their representations in document feature space:

$$\text{score}(q, d) = \frac{h(q)^\top \cdot \vec{R}_D^{(d)}}{\|h(q)\| \|\vec{R}_D^{(d)}\|}. \quad (8.3)$$

We then proceed by ranking documents $d \in D$ in decreasing order of $\text{score}(q, d)$ (Eq. 8.3) for a given query q . Note that the cosine similarity between two vectors is equivalent to their Euclidean distance if the vectors are normalized. Therefore, ranking according to Eq. 8.3 can be formulated as an (approximate) nearest neighbour search problem in a metric space.

The model proposed here, NVSM, is an extension of the LSE model [256]. NVSM/LSE are different from existing unsupervised neural retrieval models learned from scratch due to their ability to scale to collections larger than expert finding collections [257] (i.e., $\sim 10k$ entities/documents) and the assumption that words and documents are embedded in different spaces [141]. Compared to word embeddings [177, 199], NVSM/LSE learn document-specific representations instead of collection-wide representations of language; see Section 7.3.2 for a more in-depth discussion. The algorithmic contribution of NVSM over LSE comes from improvements in the objective that we introduce in the next section.

8.3.2 The objective and its optimization

We learn representations of words and documents using mini-batches of m n -gram/document pairs such that an n -gram representation—composed out of word representations—is projected nearby the document that contains the n -gram, similar to LSE [256]. The word and n -gram representations, and the transformation between them are all learned simultaneously. This is in contrast to representing documents by a weighted sum of pre-trained representations of the words that it contains [270]. A mini-batch B is constructed as follows: (1) Stochastically sample document $d \in D$ according to $P(D)$. In this chapter, we assume $P(D)$ to be uniform, similar to [286]. Note that $P(D)$ can be used to incorporate importance-based information (e.g., document length). (2) Sample a phrase of n contiguous words w_1, \dots, w_n from document d . (3) Add the phrase-document pair $(w_1, \dots, w_n; d)$ to the batch. (4) Repeat until the batch is full.

Given a batch B , we proceed by constructing a differentiable, non-convex loss function L of the parameters θ (e.g., R_V, R_D) and the parameter estimation problem is cast as an optimization problem that we approximate using stochastic gradient descent.

Denote B_i as the i -th n-gram/document pair of batch B and $B_i^{(p)}$ ($B_i^{(d)}$, respectively) as the n-gram (document, respectively) of pair B_i . Further, we introduce an auxiliary function that L2-normalizes a vector of arbitrary dimensionality:

$$\text{norm}(\vec{x}) = \frac{\vec{x}}{\|\vec{x}\|}.$$

For an n-gram/document pair B_i , the non-standardized projection of the n-gram into the k_d -dimensional document feature space is as follows:

$$\tilde{T}(B_i^{(p)}) = (f \circ \text{norm} \circ g)(B_i^{(p)}). \quad (8.4)$$

A few quick comments are in order. The function g (see Eq. 8.2) constructs an n-gram representation by averaging the word representations (embedded in the R_V parameter matrix). This allows the model to learn semantic relations between words for the purpose of **semantic matching**. That is, the model does not learn from individual words, but instead it learns from an unordered sequence (order is not preserved as we sum in Eq. 8.2) of words that constitute meaning. As documents contain multiple n-grams and n-grams are made up of multiple words, semantic similarity between words and documents is learned. In addition, the composition function g in combination with L2-normalization $\text{norm}(\cdot)$ causes words to compete in order to contribute to the resulting n-gram representation. Given that we will optimize the n -gram representation to be close to the corresponding document (as we will explain below), words that are discriminative for the document in question will learn to contribute more to the n-gram representation (due to their discriminative power), and consequently, the L2-norm of the representations of discriminative words will be larger than the L2-norm of non-discriminative words. This incorporates a notion of **term specificity** into our model.

We then estimate the per-feature sample mean and variance

$$\hat{\mathbb{E}}[\tilde{T}(B_i^{(p)})] \text{ and } \hat{\mathbb{V}}[\tilde{T}(B_i^{(p)})]$$

over batch B . The standardized projection of n-gram/document pair B_i can then be obtained as follows:

$$T(B_i^{(p)}) = \text{hard-tanh}\left(\frac{\tilde{T}(B_i^{(p)}) - \hat{\mathbb{E}}[\tilde{T}(B_i^{(p)})]}{\sqrt{\hat{\mathbb{V}}[\tilde{T}(B_i^{(p)})]}} + \beta\right), \quad (8.5)$$

where β is a k_d -dimensional bias vector parameter that captures document-independent regularities corresponding to word frequency. While vector β is learned during training, it is ignored during prediction (i.e., a nuisance parameter) similar to the position bias in click models [55] and score bias in learning-to-rank [125]. The standardization

operation reduces the internal covariate shift [123]. That is, it avoids the complications introduced by changes in the distribution of document feature vectors during learning. In addition, as the document feature vectors are learned from scratch as well, the standardization forces the learned feature vectors to be centered around the null vector. Afterwards, we apply the hard-saturating nonlinearity hard-tanh [101] such that the feature activations are between -1 and 1 .

The objective is to maximize the similarity between

$$T\left(B_i^{(p)}\right) \text{ and } \vec{R}_D^{(B_i^{(d)})},$$

while minimizing the similarity between $T\left(B_i^{(p)}\right)$ and the representations of other documents. Therefore, a document is characterized by the concepts it contains, and consequently, documents describing similar concepts will cluster together, as postulated by the **clustering hypothesis** [265]. NVSM strongly relies on the clustering hypothesis as ranking is performed according to nearest neighbour retrieval (Eq. 8.3).

Considering the full set of documents D is often costly as $|D|$ can be large. Therefore, we apply an adjusted-for-bias variant of negative sampling [106, 176, 256], where we uniformly sample negative examples from D . Adopted from [256], we define

$$P\left(\mathcal{S} \mid d, B_i^{(p)}\right) = \sigma\left(\vec{R}_D^{(d)} \cdot T\left(B_i^{(p)}\right)\right) \quad (8.6)$$

as the similarity of two representations in latent vector space, where

$$\sigma(t) = \frac{1}{1 + \exp(-t)}$$

denotes the sigmoid function and \mathcal{S} is an indicator binary random variable that says whether the representation of document d is similar to the projection of n-gram $B_i^{(p)}$.

The probability of document $B_i^{(d)}$ given phrase $B_i^{(p)}$ is then approximated by uniformly sampling z contrastive examples:

$$\begin{aligned} \log \tilde{P}\left(B_i^{(d)} \mid B_i^{(p)}\right) = & \quad (8.7) \\ & \frac{z+1}{2z} \left(z \log P\left(\mathcal{S} \mid B_i^{(d)}, B_i^{(p)}\right) + \right. \\ & \left. \sum_{\substack{k=1, \\ d_k \sim U(D)}}^z \log \left(1.0 - P\left(\mathcal{S} \mid d_k, B_i^{(p)}\right) \right) \right), \end{aligned}$$

where $U(D)$ denotes the uniform distribution over documents D , the distribution used for obtaining negative examples [106, 256]. Then, the **loss function** we use to optimize

our model is Eq. 8.7 averaged over the instances in batch B :

$$\begin{aligned} L(R_V, R_D, W, \beta \mid B) \\ = \frac{1}{m} \sum_{i=1}^m \log \tilde{P}(B_i^{(d)} \mid B_i^{(p)}) \\ + \frac{\lambda}{2m} \left(\sum_{i,j} R_{V,i,j}^2 + \sum_{i,j} R_{D,i,j}^2 + \sum_{i,j} W_{i,j}^2 \right), \end{aligned} \quad (8.8)$$

where λ is a weight regularization hyper-parameter. We optimize our parameters θ (R_V , R_D , W and β) using Adam [133], a first-order gradient-based optimization function for stochastic objective functions that is very similar to momentum. The update rule for parameter θ given a batch B at batch update time t equals:

$$\theta^{(t+1)} \leftarrow \theta^{(t)} - \alpha \frac{\hat{m}_\theta^{(t)}}{\sqrt{\hat{v}_\theta^{(t)} + \epsilon}}, \quad (8.9)$$

where $\hat{m}_\theta^{(t)}$ and $\hat{v}_\theta^{(t)}$, respectively, are the first and second moment estimates (over batch update times) [133] of the gradient of the loss $\frac{\partial L}{\partial \theta}(R_V, R_D, W, \beta \mid B)$ w.r.t. parameter θ at batch update time t and $\epsilon = 10^{-8}$ is a constant to ensure numerical stability. The use of this optimization method causes every parameter to be updated with every batch, unlike regular stochastic gradient descent, where the only parameters that are updated are those for which there is a non-zero gradient of the loss. This is important in NVSM due to the large number of word and document vectors.

The algorithmic contributions of NVSM over the LSE model that we introduced in the previous chapter are the components of the objective mentioned next. Eq. 8.4 forces individual words to compete in order to contribute to the resulting n-gram representation. Consequently, non-discriminative words will have a small L2-norm. In Eq. 8.5 we perform standardization to reduce the internal covariate shift [123]. In addition, the standardization forces n-gram representations to distinguish themselves only in the dimensions that matter for matching. Frequent words are naturally prevalent in n-grams, however, they have low discriminative power as they are non-specific. The bias β captures word frequency regularities that are non-discriminative for the semantic concepts within the respective n-gram/document pairs and allows the transformation in Eq. 8.1 to focus on concept matching. The re-weighting of the positive instance in Eq. 8.7 removes a dependence on the number of negative examples z where a large z presented the model with a bias towards negative examples.

8.3.3 Implementation

The major cause of technical challenges of the NVSM training procedure is not due to time complexity, but rather space restrictions. This is because we mitigate expensive computation by estimating vector space models using graphics processing units (GPUs). The main limitation of these massively-parallel computation devices is that they rely on their own memory units. Consequently, parameters and intermediate results of the

training procedure need to persist in the GPU memory space. The asymptotic space complexity of the parameters equals:

$$O \left(\underbrace{|V| \times k_w}_{\text{word representations } R_V} + \underbrace{k_d \times k_w}_{\text{transform } W} + \underbrace{|D| \times k_d}_{\text{document representations } R_D} \right).$$

In addition, Eq. 8.9 requires us to keep the first and second moment of the gradient over time for every parameter in memory. Therefore, for every parameter, we retain three floating point values in memory at all times. For example, if we have a collection of 1M documents (256-dim.) with a vocabulary of 64K terms (300-dim.), then the model has $\sim 275\text{M}$ parameters. Consequently, under the assumption of 32-bit floating point, the resident memory set has a best-case least upper bound of 3.30GB memory. The scalability of our method—similar as with all latent vector space models—is determined by the number of documents within the collection. However, the current generation of GPUs—that typically boast around 12GB of memory—can comfortably be used to train models of collections consisting of up to 2 million documents. In fact, next-generation GPUs have double the amount of memory—24GB—and this amount will likely increase with the introduction of future processing units [276]. This, and the development of distributed GPU technology [2], leads us to believe that the applicability of our approach to larger retrieval domains is simply a matter of time [186].

In addition to the scarcity of memory on current generation GPUs, operations such as the averaging of word to n-gram representations (Eq. 8.2) can be performed in-place. However, these critical optimizations are not available in general-purpose machine learning toolkits. Therefore, we have implemented the NVSM training procedure directly in C++/CUDA, such that we can make efficient use of sparseness and avoid unnecessary memory usage. In addition, models are stored in the open HDF5 format [86] and the toolkit provides a Python module that can be used to query trained NVSM models on the CPU. This way, a trained NVSM can easily be integrated in existing applications. The toolkit is licensed under the permissive MIT open-source license.¹

8.4 Experimental setup

8.4.1 Research questions

In this chapter we investigate the viability of neural representation learning methods for semantic matching in document search. As indicated in the introduction of this chapter, we seek to answer the following research questions:

RQ4.1 How does NVSM compare to other latent vector space models, such as doc2vec [141], word2vec [177, 270], LSI [70], LDA [38] and LSE [256], on the document retrieval task?

In particular, how does it compare to other methods that represent queries/documents as low-dimensional vectors? What is the difference in performance with purely lexical

¹<https://github.com/cvangysel/cuNVSM>

models that perform exact term matching and represent queries/documents as bag-of-words vectors?

RQ4.2 For what proportion of queries does NVSM perform better than the other rankers?

Does NVSM improve over other retrieval models only on a handful of queries? Instead of computing averages over queries, what if we look at the pairwise differences between rankers?

RQ4.3 What gains does NVSM bring when combined with a lexical QLM and a competing state-of-the-art vector space model?

Can we use the differences that we observe in **RQ4.2** between NVSM, QLM and other latent vector space models to our advantage to improve retrieval performance? Can we pinpoint where the improvements of NVSM come from?

RQ4.4 Do NVSMs exhibit regularities that we can link back to well-understood document collection statistics used in traditional retrieval models?

Traditional retrieval models such as generative language models are known to incorporate corpus statistics regarding term specificity [215, 237] and document length [286]. Are similar corpus statistics present in NVSM and what does this tell us about the ranking task?

8.4.2 Benchmark datasets & experiments

In this chapter we are interested in query/document matching. Therefore, we evaluate NVSM on newswire article collections from TREC. Other retrieval domains, such as web search or social media, deal with various aspects, such as document freshness/importance and social/hyperlink/click graphs, that may obfuscate the impact of matching queries to documents. Therefore, we follow the experimental setup of Zhai and Lafferty [286] and use four article retrieval sub-collections from the TIPSTER corpus [107]: Associated Press 88-89 (AP88-89), Financial Times (FT), LA Times (LA) and Wall Street Journal (WSJ) [108]. In addition, we consider the Robust04 collection that constitutes of Disk 4/5 of the TIPSTER corpus without the Congressional Record and the New York Times collection that consists of articles written and published by the New York Times between 1987 and 2007. For evaluation, we take topics 50–200 from TREC 1–3² (AP88-89, WSJ), topics 301–450 from TREC 6–8² (FT, LA) [244], topics 301–450, 601–700 from Robust04 [269] and the 50 topics assessed by NIST (a subset of the Robust04 topics judged for the NY collection) during the TREC 2017 Common Core Track [8]. From each topic we take its title as corresponding query. Topics without relevant documents are filtered out. We randomly create a split³ of validation (20%) and test (80%) queries (with the exception of the NY collection). For the NY collection, we select the hyperparameter configuration that optimizes the Robust04 validation set on the Robust04 collection and we take the 50 queries assessed by NIST and their judgments—specifically created for the NY collection—as our test

²We only consider judgments corresponding to each of the sub-collections.

³The validation/test splits can be found at <https://github.com/cvangysel/cuNVSM>.

set. This way, method hyperparameters are optimized on the validation set (as described in Section 8.4.3) and the retrieval performance is reported on the test set; see Table 8.1.

The inclusion of early TREC collections (AP88-89, FT, LA and WSJ) is motivated by the fact that during the first few years of TREC, there was a big emphasis on submissions where the query was constructed manually from each topic and interactive feedback was used [108]. That is, domain experts repeatedly formulated manual queries using the full topic (title, description, narrative), observed the obtained rankings and then reformulated their query in order to obtain better rankings. Consequently, these test collections are very useful when evaluating semantic matches as relevant documents do not necessarily contain topic title terms. From TREC-5 and onwards, less emphasis was put on rankings generated using interactive feedback and shifted towards automated systems only [109, footnote 1]. In fact, within the Robust04 track, only automated systems were submitted [269, Section 2] due to the large number of (a) documents ($\sim 500K$) and (b) topics (250). In the 2017 Common Core Track [8], interactive rankings were once again submitted by participants, in addition to rankings obtained by the latent vector space model presented in this chapter.

We address **RQ4.1** by comparing NVSM to latent retrieval models (detailed in Section 8.4.3). In addition, we perform a per-query pairwise comparison of methods where we look at what method performs best for each query in terms of Mean Average Precision (MAP)@1000 (**RQ4.2**). A method performs better or worse than another if the absolute difference in MAP@1000 exceeds $\delta = 0.01$; otherwise, the two methods perform similar. To address **RQ4.3**, we consider the combinations (Section 8.4.3.2) of QLM with NVSM and the strongest latent vector space baseline of **RQ4.1**. That is, word2vec where the summands are weighted using self-information. In addition, we look at the correlation between per-query `titlestat_rel` (see Section 8.4.4) and the pairwise differences in MAP@1000 between NVSM and all the other retrieval models. A positive correlation indicates that NVSM is better at lexical matching than the other method, and vice versa for a negative correlation. For **RQ4.4**, we examine the relation between the collection frequency CF_w and the L2-norm of their word embeddings $g(w)$ for all terms $w \in V$.

8.4.3 Retrieval models considered for comparison

The document collection is first indexed by Indri⁴ [238]. Retrieval models not implemented by Indri access the underlying tokenized document collection using `pyndri` [262]. This way, all methods compared in this chapter parse the text collection consistently.

8.4.3.1 Models compared

The key focus of this chapter is the alleviation of the vocabulary gap in information retrieval and consequently, in theory, we score all documents in each collection for every query. In practice, however, we rely on nearest neighbor search algorithms to retrieve the top-k documents [43]. Note that this is in contrast to many other semantic matching methods [4, 190, 290] that have only been shown to perform well in document re-ranking

⁴Stopwords are removed using the standard stopword list of Indri.

Table 8.1: Overview of the retrieval benchmarks. T and V denote the test and validation sets, respectively. Arithmetic mean and standard deviation are reported wherever applicable.

	AP88-89	FT	LA
Collection (training)			
Documents	164,597	210,158	131,896
Document length	461.63 ± 243.61	399.68 ± 366.41	502.18 ± 519.58
Unique terms	2.67×10^5	3.05×10^5	2.67×10^5
Queries (testing)			
	(T) 119	(T) 116	(T) 113
	(V) 30	(V) 28	(V) 30
Query terms	5.06 ± 3.14	2.50 ± 0.69	2.48 ± 0.69
Relevant documents	(T) 111.45 ± 136.01	(T) 34.91 ± 42.94	(T) 24.83 ± 34.31
	(V) 86.43 ± 72.63	(V) 30.46 ± 26.97	(V) 24.30 ± 21.19
Collection (training)			
Documents	1,855,658	528,155	173,252
Document length	572.18 ± 605.82	479.72 ± 869.27	447.51 ± 454.75
Unique terms	1.35×10^6	7.83×10^5	2.50×10^5
Queries (testing)			
	(T) 50	(T) 200	(T) 120
		(V) 49	(V) 30
Query terms	6.58 ± 0.70	5.28 ± 0.74	5.05 ± 3.14
Relevant documents	(T) 180.04 ± 132.74	(T) 70.33 ± 73.68	(T) 96.99 ± 93.30
		(V) 68.27 ± 77.41	(V) 101.93 ± 117.65

scenarios where an initial pool of candidate documents is retrieved using a lexical matching method. However, candidate documents most affected by the vocabulary gap (i.e., relevant documents that do not contain any query terms) will simply remain undiscovered in a re-ranking scenario and consequently we compare NVSM only to latent vector space models that can be queried using a nearest neighbor search.

The following latent vector space models are compared:

1. doc2vec (doc2vec (d2v)) [141] with the distributed memory architecture. The pre-processing of document texts to learn latent document representations is a topic of study by itself and its effects are outside the scope of this work. Consequently, we disable vocabulary filtering and frequent word subsampling in order to keep the input to all representation learning algorithms consistent. We sweep the one-sided window size and the embedding size respectively in partitions $\{x/2 \mid x = 4, 6, 8, 10, 12, 16, 24, 32\}$ and $\{64, 128, 256\}$ on the validation set. Models are trained for 15 iterations on the validation set and we select the model iteration that performs best on the validation set. Documents are ranked in decreasing order of the cosine similarity between the document representation and the average of the word embeddings in the query.
2. word2vec (word2vec (w2v)) [177, 270] with the Skip-Gram architecture. We follow the method introduced by Vulić and Moens [270] where query/document representations are constructed by composing the representations of the words contained within them. We consider both the unweighted sum (`add`) and the sum of vectors weighted by the term's self-information (`self-information (si)`). Self-information is a term specificity measure similar to Inverse Document Frequency (IDF) [61]. The hyperparameters of word2vec are swept in the same manner as doc2vec.
3. Latent Semantic Indexing (LSI) [70] with TF-IDF weighting and the number of topics $K \in \{64, 128, 256\}$ optimized on the validation set.
4. Latent Dirichlet Allocation (LDA) [38] with $\alpha = \beta = 0.1$ and the number of topics $K \in \{64, 128, 256\}$ optimized on the validation set. We train the model for 100 iterations or until topic convergence is achieved. Documents are ranked in decreasing order of the cosine similarity between the query topic distribution and the document topic distribution.
5. Representation learning methods LSE [256] and NVSM (this chapter). For hyperparameters, we largely follow the findings of [256]: word representation dim. $k_w = 300$, number of negative examples $z = 10$, learning rate $\alpha = 0.001$, regularization lambda $\lambda = 0.01$. For LSE, batch size $m = 4096$ (as in [256]), while for NVSM the batch size $m = 51200$ (empirically determined on a holdout document collection that we did not include in this chapter). The dimensionality of the document representations $k_d \in \{64, 128, 256\}$ and the n-gram size $n \in \{4, 6, 8, 10, 12, 16, 24, 32\}$ are optimized on the validation set. Similar to d2v, models are trained for 15 iterations on the training set and we select the model iteration that performs best on the validation set; a single iteration consists of $\lceil \frac{1}{m} \sum_{d \in D} (|d| - n + 1) \rceil$ batches.

In addition, we consider lexical language models (QLM) [286] using the Indri engine with both Dirichlet (Dirichlet (d)) and Jelinek-Mercer (Jelinek-Mercer (jm)) smoothing; smoothing hyperparameters $\mu \in \{125, 250, 500, 750, 1000, 2000, 3000, 4000, 5000\}$ and $\lambda \in \{x \mid k \in \mathbb{N}_{>0}, k \leq 20, x = k/20\}$, respectively, are optimized on the validation set. The retrieval effectiveness of QLM is provided as a point of reference in **RQ4.1**. For **RQ4.2**, the QLM is used as a lexical retrieval model that is fused with latent vector space models to provide a mixture of lexical and semantic matching.

For the latent vector spaces (d2v, LSI, LDA, LSE and NVSM), the vocabulary size is limited to the top-60k most frequent words as per [256], given that latent methods rely on word co-occurrence in order to learn latent relations. For doc2vec, word2vec, LSI and LDA we use the Gensim⁵ implementation; the neural representation learning methods use our open source CUDA implementation described in Section 8.3.3; see footnote 1.

8.4.3.2 Combinations of QLM and latent features

We combine individual rankers by performing a grid search on the weights of a linear combination using 20-fold cross validation on the test sets (Table 8.1). For QLM, feature values are the log-probabilities of the query given the document, while for the latent features (NVSM and w2v-si), we use the cosine similarity between query/document representations. For every feature weight, we sweep between 0.0 and 1.0 with increments of 0.0125 on the fold training set. Individual features are normalized per query such that their values lie between 0 and 1. We select the weight configuration that achieves highest Mean Average Precision on the training set and use that configuration to score the test set. During scoring of the fold test set, we take the pool of the top-1k documents ranked by the individual features as candidate set.

8.4.4 Evaluation measures and statistical significance

To address **RQ4.1**, **RQ4.2** and **RQ4.3**, we report Mean Average Precision at rank 1000 (MAP@1000), Normalized Discounted Cumulative Gain at rank 100 (NDCG@100) and Precision at rank 10 (Precision (P)@10) to measure retrieval effectiveness. For **RQ4.3**, we also look at the per-query `titlestat_rel` [45], the expected normalized term overlap between query and document. All evaluation measures are computed using TREC’s official evaluation tool, `trec_eval`.⁶ Wherever reported, significance of observed differences is determined using a two-tailed paired Student’s t-test [234] (** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$). For correlation coefficients, significance is determined using a permutation test ([†] $p < 0.01$). For **RQ4.4**, we use Welch’s t-test to determine whether the mean L2-norm of mid-frequency (middle-50%) words is significantly different from the mean L2-norm of low- (bottom 25%) and high-frequency (top 25%) words.

⁵<https://github.com/RaRe-Technologies/gensim>

⁶https://github.com/usnistgov/trec_eval

8.5 Results

First, we present a comparison between methods (**RQ4.1**) on ad-hoc document retrieval, followed by a per-query comparison between methods (**RQ4.2**) and a combination experiment where we combine latent features with the lexical QLM (**RQ4.3**). We then relate regularities learned by the model to traditional retrieval statistics (**RQ4.4**).

8.5.1 Performance of NVSM

RQ4.1 Table 8.2 shows the retrieval results for ad-hoc document retrieval on the newswire article collections (Section 8.4.2).

We see that NVSM outperforms all other latent rankers on all benchmarks. In particular, NVSM significantly outperforms (MAP@1000) the word2vec-based method that weighs word vectors according to self-information (significance is not achieved on NY). This is an interesting observation as NVSM is trained from scratch without the use of hand-engineered features (i.e., self-information). Compared to the lexical QLM, NVSM performs better on the AP88-89 and WSJ benchmarks. However, it is known that no single ranker performs best on all test sets [151, 231]. In addition, NVSM is a latent model that performs a different type of matching than lexical models. Therefore, we first examine the per-query differences between rankers (**RQ4.2**) and later we will examine the complementary nature of the two types of matching by evaluating combinations of different ranking features (**RQ4.3**).

8.5.2 Query-level analysis

RQ4.2 Fig. 8.1 shows the distribution of queries where one individual ranker performs better than the other ($|\Delta\text{MAP@1000}| > \delta$).

We observe similar trends across all benchmarks where NVSM performs best compared to all latent rankers. One competing vector space model, w2v-si, stands out as it is the strongest baseline and performs better than NVSM on 20 to 35% of queries, however, NVSM still beats w2v-si overall and specifically on 40 to 55% of queries. Moreover, Fig. 8.1 shows us that QLM and NVSM make very different errors. This implies that the combination of QLM, w2v-si and NVSM might improve performance even further.

While answering **RQ4.1**, we saw that for some benchmarks, latent methods (i.e., NVSM) perform better than lexical methods. While the amount of semantic matching needed depends on various factors, such as the query intent being informational or navigational/transactional [44], we do see in Fig. 8.1 that NVSM performs considerably better amongst latent methods in cases where latent methods perform poorly (e.g., Robust04 in Fig. 8.1e). Can we shine more light on the difference between NVSM and existing latent methods? We answer this question in the second half of the next section.

8.5.3 Semantic vs. lexical matching

RQ4.3 Beyond individual rankers, we now also consider combinations of the two best-performing vector space models with QLM [231] (see Section 8.4.3.2 for

Table 8.2: Comparison of NVSM with lexical (QLM with Dirichlet and Jelinek-Mercer smoothing) and latent (d2v, w2v, LSI, LDA and LSE) retrieval models (Section 8.4.3) on article search benchmarks (Section 8.4.2). Significance (Section 8.4.4) is computed between w2v-si and NVSM. Bold values indicate the highest measure value for latent features.

	AP88-89			FT			LA		
	MAP	NDCG	P@10	MAP	NDCG	P@10	MAP	NDCG	P@10
Bag-of-words features (for reference)									
QLM (jm)	0.199	0.346	0.365	0.218	0.356	0.283	0.182	0.331	0.221
QLM (d)	0.216	0.370	0.392	0.240	0.381	0.296	0.198	0.348	0.239
Latent features (for comparison)									
d2v	0.002	0.011	0.009	0.001	0.004	0.004	0.006	0.020	0.014
LDA	0.039	0.077	0.078	0.009	0.028	0.013	0.004	0.015	0.010
LSI	0.131	0.228	0.242	0.029	0.068	0.052	0.033	0.079	0.061
LSE	0.144	0.281	0.299	0.060	0.140	0.123	0.067	0.152	0.105
w2v-add	0.216	0.370	0.393	0.125	0.230	0.195	0.105	0.212	0.159
w2v-si	0.230	0.383	0.418	0.141	0.250	0.204	0.131	0.242	0.179
NVSM	0.257**	0.418**	0.425	0.172**	0.302***	0.239*	0.166**	0.300***	0.209*
	NY			Robust04			WSJ		
	MAP	NDCG	P@10	MAP	NDCG	P@10	MAP	NDCG	P@10
Bag-of-words features (for reference)									
QLM (jm)	0.158	0.270	0.376	0.201	0.359	0.369	0.175	0.315	0.345
QLM (d)	0.188	0.318	0.486	0.224	0.388	0.415	0.204	0.351	0.398
Latent features (for comparison)									
d2v	0.081	0.179	0.272	0.068	0.177	0.194	0.003	0.015	0.011
LDA	0.009	0.027	0.022	0.003	0.010	0.009	0.038	0.082	0.076
LSI	0.020	0.041	0.040	0.022	0.059	0.060	0.101	0.181	0.207
LSE	0.000	0.000	0.000	0.013	0.050	0.054	0.098	0.208	0.245
w2v-add	0.081	0.160	0.216	0.075	0.177	0.194	0.175	0.322	0.372
w2v-si	0.092	0.173	0.220	0.093	0.208	0.234	0.185	0.330	0.391
NVSM	0.117	0.208	0.296*	0.150***	0.287***	0.298***	0.208**	0.351	0.370




Figure 8.1: Per-query pairwise ranker comparison between NVSM and the QLM (d), QLM (jm), d2v, w2v, LSI, LDA and LSE rankers. For every bar, the dotted/green area, solid/orange and red/slashed areas respectively depict the portion of queries for which NVSM outperforms, ties or loses against the other ranker. One ranker outperforms the other if the absolute difference in MAP@1000 between both rankers exceeds δ .

Table 8.3: Evaluation of latent features (d2v, w2v, LSI, LDA, LSE and NVSM) as a complementary signal to QLM (d) (Section 8.4.3.2). Significance (Section 8.4.4) is computed between QLM (d) + w2v-si and the best performing combination.

	AP88-89		
	MAP	NDCG	P@10
QLM (d)	0.216	0.370	0.392
QLM (d) + w2v-si	0.279 (+29%)	0.437 (+18%)	0.450 (+14%)
QLM (d) + NVSM	0.289 (+33%)	0.444 (+20%)	0.473 (+20%)
QLM (d) + w2v-si + NVSM	0.307*** (+42%)	0.466*** (+26%)	0.498*** (+27%)
	FT		
	MAP	NDCG	P@10
QLM (d)	0.240	0.381	0.296
QLM (d) + w2v-si	0.251 (+4%)	0.393 (+3%)	0.313 (+6%)
QLM (d) + NVSM	0.251 (+4%)	0.401 (+5%)	0.322 (+9%)
QLM (d) + w2v-si + NVSM	0.258 (+7%)	0.406 (+6%)	0.322 (+9%)
	LA		
	MAP	NDCG	P@10
QLM (d)	0.198	0.348	0.239
QLM (d) + w2v-si	0.212 (+7%)	0.360 (+3%)	0.236 (-1%)
QLM (d) + NVSM	0.220 (+11%)	0.376 (+7%)	0.244 (+1%)
QLM (d) + w2v-si + NVSM	0.226*** (+14%)	0.378*** (+8%)	0.250** (+4%)
	NY		
	MAP	NDCG	P@10
QLM (d)	0.188	0.318	0.486
QLM (d) + w2v-si	0.206 (+9%)	0.333 (+4%)	0.494 (+1%)
QLM (d) + NVSM	0.222 (+18%)	0.355** (+11%)	0.520 (+6%)
QLM (d) + w2v-si + NVSM	0.222*** (+18%)	0.353 (+10%)	0.526** (+8%)
	Robust04		
	MAP	NDCG	P@10
QLM (d)	0.224	0.388	0.415
QLM (d) + w2v-si	0.232 (+3%)	0.399 (+2%)	0.428 (+2%)
QLM (d) + NVSM	0.247 (+10%)	0.411 (+6%)	0.448*** (+7%)
QLM (d) + w2v-si + NVSM	0.247*** (+10%)	0.412*** (+6%)	0.446 (+7%)
	WSJ		
	MAP	NDCG	P@10
QLM (d)	0.204	0.351	0.398
QLM (d) + w2v-si	0.254 (+24%)	0.410 (+16%)	0.454 (+13%)
QLM (d) + NVSM	0.248 (+21%)	0.396 (+12%)	0.425 (+6%)
QLM (d) + w2v-si + NVSM	0.271*** (+32%)	0.426*** (+21%)	0.456 (+14%)

details) in Table 8.3.

If we consider the QLM paired with either w2v-si or NVSM, we see that the combination involving NVSM outperforms the combination with w2v-si on four out of six benchmarks (AP88-89, LA, NY, Robust04). However, Figure 8.1 shows that NVSM and

Table 8.4: Correlation coefficients between `titlestat_rel` and $\Delta_{\text{MAP@1000}}$ between NVSM and the other methods. A positive correlation indicates that NVSM is better at lexical matching, while a negative correlation indicates that NVSM is worse at lexical matching than the alternative. Significance (Section 8.4.4) is computed using a permutation test.

	AP88-89	FT	LA	NY	Robust04	WSJ
Bag-of-words features						
QLM (jm)	-0.102	-0.355 [†]	-0.188 [†]	-0.375 [†]	-0.337 [†]	-0.168 [†]
QLM (d)	-0.211 [†]	-0.456 [†]	-0.206 [†]	-0.432 [†]	-0.374 [†]	-0.275 [†]
Latent features						
d2v	0.625 [†]	0.415 [†]	0.473 [†]	0.195	0.161 [†]	0.459 [†]
LDA	0.545 [†]	0.406 [†]	0.497 [†]	0.380 [†]	0.309 [†]	0.376 [†]
LSI	0.420 [†]	0.376 [†]	0.528 [†]	0.344 [†]	0.217 [†]	0.282 [†]
LSE	0.506 [†]	0.361 [†]	0.396 [†]	0.374 [†]	0.289 [†]	0.270 [†]
w2v-add	0.374 [†]	0.275 [†]	0.450 [†]	0.283 [†]	0.230 [†]	0.393 [†]
w2v-si	0.232 [†]	0.171 [†]	0.316 [†]	0.203	0.247 [†]	0.357 [†]

w2v-si outperform each other on different queries as well. Can we use this difference to our advantage?

The addition of NVSM to the QLM + w2v-si combination yields an improvement in terms of MAP@1000 on all benchmarks. Significance is achieved in five out of six benchmarks. In the case of NY and Robust04, the combination of all three rankers (QLM + w2v-si + NVSM) performs at about the same level as the combination of QLM + NVSM. However, the addition of NVSM to the QLM + w2v-si combination still creates a significant improvement over just the combination involving QLM and w2v only. For FT, the only benchmark where no significance is achieved, we do see that the relative increase in performance nearly doubles from the addition of NVSM. Consequently, we can conclude that the NVSM adds an additional matching signal.

Let us return to the question raised at the end of the previous section (Section 8.5.2): what exactly does the NVSM add in terms of content matching? We investigate this question by determining the amount of semantic matching needed. For each query, we compute `titlestat_rel` (Section 8.4.4), the expected normalized overlap between query terms and the terms of relevant document. If `titlestat_rel` is close to 1.0 for a particular query, then the query requires mostly lexical matching; on the other hand, if `titlestat_rel` is near 0.0 for a query, then none of the query’s relevant document contain the query terms and semantic matching is needed. We continue by examining the per-query pairwise difference ($\Delta_{\text{MAP@1000}}$) between NVSM and the remaining lexical (QLM) and latent (d2v, LDA, LSI, LSE) features. If $\Delta_{\text{MAP@1000}} > 0$, then NVSM performs better than the other method and vice versa if $\Delta_{\text{MAP@1000}} < 0$. Table 8.4 shows the Pearson correlation between `titlestat_rel` and $\Delta_{\text{MAP@1000}}$. A positive correlation, as is the case for d2v, w2v, LDA, LSI and LSE, indicates that NVSM performs better on queries that require lexical matching. Conversely, a negative correlation, such as observed for both variants of QLM, indicates that QLM performs better on queries that require lexical matching than NVSM. Combining this observation with the conclusion to **RQ4.2** (i.e., NVSM generally improves upon latent methods), we conclude that, in addition to semantic matching, NVSM also performs well in cases

where lexical matching is needed and thus contributes a hybrid matching signal.

8.5.4 NVSM and Luhn significance

If NVSM performs better at lexical matching than other latent vector space models, does it then also contain regularities associated with term specificity?

RQ4.4 Fig. 8.2 shows the L2-norm of individual term representations for LSE (left scatter plots) and NVSM (right scatter plots).

Luhn [154] measures the significance of words based on their frequency. They specify a lower and upper frequency cutoff to exclude frequent and infrequent words. For NVSM (scatter plot on the right for every benchmark), we find that infrequent and frequent terms have a statistically significant ($p < 0.01$) smaller L2-norm than terms of medium frequency (Section 8.4.4). This observation is further motivated by the shape of the relation between collection frequency in the collection and the L2-norm of term representations in Fig. 8.2. The key observation that—within NVSM representations—terms of medium frequency are of greater importance (i.e., higher L2-norm) than low- or high-frequency terms closely corresponds to the theory of Luhn significance. Particularly noteworthy is the fact that the NVSM learned this relationship from an unsupervised objective directly, without any notion of relevance. The scatter plots on the left for every benchmark in Fig. 8.2 shows the same analysis for LSE term representations. Unlike with NVSM, we observe that the L2-norm of term representations grows linearly with the term collection frequency, and consequently, high-frequency terms are of greater importance within LSE representations. Therefore, the key difference between NVSM and LSE is that NVSM learns to better encode term specificity.

8.6 Discussion and analysis

In this section, we investigate the impact of the judgement bias (Section 8.6.1). We then proceed by giving guidelines on how to deploy NVSM in the absence of a validation set (Section 8.6.2).

8.6.1 An investigation of judgement bias

In what capacity does the judgement bias inherent to the construction of TREC test collections affect the evaluation of novel retrieval models such as NVSM? The relevance assessments of TREC test collections are created using a pooling strategy of the rankings produced by participants of the TREC ad-hoc track [244]. This strategy is known to introduce a pooling bias [45, 112, 150]: rankings that have fewer documents judged may be scored lower just because of this fact. Does judgement bias influence the evaluation of NVSM compared to lexical language models?

Fig. 8.3 shows the average number of judged documents at different ranks for bag-of-words methods (QLM) and the NVSM on the article search benchmarks (Section 8.4.2). Here we omit the NY benchmark, as NVSM-based rankings were used during the construction of the NY test collection. We can see that for AP88-89 and WSJ, all

methods have approximately the same number of judged documents over all rank cut-offs. However, on WSJ, QLM with Dirichlet smoothing has more judged documents for lower rank cut-offs (cut-off 1, 5, 10 and 20) and therefore we can conclude that there is a bias in the evaluation towards this method.

Interestingly, on the AP88-89 and WSJ benchmarks, as the rank cut-off increases from 1–50, we see that the NVSM retrieves more, or the same number of, judged documents as the lexical methods in Fig. 8.3a–8.3e. This indicates that the gains obtained by NVSM are likely to originate from the torso of the ranking instead of the top. The judgement bias is even more prevalent on the FT and LA benchmarks. There, we observe a difference of about 10% fewer judged documents, consistent across all rank cut-offs, for NVSM.

We conclude this analysis as follows. There is an indication of a judgement bias against NVSM and more judgements are required in order to determine the true ranking amongst methods. In addition, ignoring unjudged documents does not solve the problem, as newer methods (i.e., NVSM) are more likely to retrieve relevant unjudged documents than the methods used to construct the original pools (e.g., QLM).

8.6.2 Unsupervised deployment

In our experiments, we use a validation set for model selection (training iterations and hyperparameters). However, in many cases relevance labels are unavailable. Fortunately, NVSM learns representations of the document collection directly and does not require query-document relevance information. How can we choose values for the hyperparameters of NVSM in the absence of a validation set? For the majority of hyperparameters (Section 8.4.3) we follow the setup of previous work [256, 257]. We are, however, still tasked with the problem of choosing (a) the number of training iterations, (b) the dimensionality of the document representations k_d , and (c) the size of the n -grams used for training. We choose the dimensionality of the document representations $k_d = 256$ as the value was reported to work well for LSE [256]. Fig. 8.4 shows that MAP@1000 converges as the number of training iterations increases for different n -gram widths. Therefore, we train NVSM for 15 iterations and select the last iteration model.

The final remaining question is the choice of n -gram size used during training. This parameter has a big influence on model performance as it determines the amount of context from which semantic relationships are learned. Therefore, we propose to combine different vector spaces trained using different n -gram widths as follows. We write \mathcal{N} for the set of all k for which we construct an NVSM using k -grams. For a given query q , we rank documents $d \in D$ in descending order of:

$$\text{score}_{\text{ensemble}}(q, d) = \sum_{k \in \mathcal{N}} \frac{\text{score}_{k\text{-grams}}(q, d) - \mu_{k\text{-grams}, q}}{\sigma_{k\text{-grams}, q}}, \quad (8.10)$$

where $\text{score}_{k\text{-grams}}(q, d)$ is Eq. 8.3 for NVSM of k -grams and

$$\begin{aligned} \mu_{k\text{-grams}, q} &= \hat{\mathbb{E}}[\text{score}_{k\text{-grams}}(q, D)] \\ \sigma_{k\text{-grams}, q} &= \sqrt{\hat{\mathbb{V}}[\text{score}_{k\text{-grams}}(q, D)]}, \end{aligned} \quad (8.11)$$

Table 8.5: Comparison with single cross-validated NVSM and ensemble of NVSM through the unsupervised combination of models trained on differently-sized n -grams ($\mathcal{N} = \{2, 4, 8, 10, 12, 16, 24, 32\}$). Significance (Section 8.4.4) is computed between NVSM and the ensemble of NVSM.

	AP88-89			FT		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.257	0.416	0.424	0.170	0.299	0.236
8 NVSMs (ensemble)	0.282 ***	0.453 ***	0.466 ***	0.212 ***	0.352 ***	0.282 ***
	LA			NY		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.168	0.302	0.211	0.117	0.208	0.296
8 NVSMs (ensemble)	0.188 ***	0.331 ***	0.230 ***	0.121	0.238 ***	0.324 *
	Robust04			WSJ		
	MAP	NDCG	P@10	MAP	NDCG	P@10
1 NVSM (cross-validated)	0.150	0.287	0.297	0.208	0.351	0.372
8 NVSMs (ensemble)	0.171 ***	0.323 ***	0.331 ***	0.225 ***	0.385 ***	0.423 ***

denote the sample expectation and sample variance over documents D that are estimated on the top-1000 documents returned by the individual models, respectively. That is, we rank documents according to the sum of the standardized scores of vector space models trained with different n -gram widths. The score aggregation in Eq. 8.10 is performed without any a priori knowledge about the n -gram sizes. Table 8.5 lists the performance of the unsupervised ensemble, where every model was trained for 15 iterations, against a single cross-validated model. We see that the unsupervised ensemble always outperforms (significantly in terms of MAP@1000 for all benchmarks except NY) the singleton model. Hence, we can easily deploy NVSM without any supervision and, surprisingly, it will perform better than individual models optimized on a validation set.

8.7 Summary

We proposed the Neural Vector Space Model (NVSM) that learns representations of a document collection in an unsupervised manner.

We showed that NVSM performs better than existing latent vector space/bag-of-words approaches. NVSM performs lexical and semantic matching in a latent space. NVSM provides a complementary signal to lexical language models. In addition, we showed that NVSM automatically learns a notion of term specificity. Finally, we gave advice on how to select values for the hyperparameters of NVSM. Interestingly, an unsupervised ensemble of multiple models trained with different hyperparameters performs better than a single cross-validated model.

The evidence that NVSM provides a notion of lexical matching tells us that latent vector space models are not limited to only semantic matching. While the framework presented in this chapter focuses on a single unsupervised objective, additional objectives (i.e., document/document or query/document similarity) can be incorporated to improve retrieval performance.

The LSE model [256]—introduced in Chapter 7—improved the learning time complexity of earlier entity retrieval models (Chapter 5) [257] such that they scale to $\sim 100k$ retrievable items (i.e., entities). However, as shown in Table 8.2, LSE performs poorly on article retrieval benchmarks. In this chapter, we extend LSE and learn vector spaces of $\sim 500k$ documents that perform better than existing latent vector spaces. As mentioned in the introduction, the main challenge for latent vector spaces is their limited scalability to large document collections due to space complexity. The observation that retrieval is not only impacted by the vector space representation of the relevant document, but also of the documents surrounding it, raises non-trivial questions regarding the distribution of document vectors over multiple machines. While there have been efforts towards distributed training of neural models, the application of distributed learning algorithms is left for future work. The unsupervised objective that learns from word sequences is limited by its inability to deal with very short documents. While this makes the unsupervised objective less applicable in domains such as web search, unsupervised bag-of-words approaches have the opposite problem of degrading performance when used to search over long documents. With respect to incremental indexing, there is currently no theoretically sound way to obtain representations for new documents

that were added to the collection after the initial estimation of a NVSM. In the case of LDA or LSI, representations for new documents can be obtained by transforming bag-of-words vectors to the latent space. However, as the LDA/LSI transformation to the latent space is not updated after estimating the LDA/LSI model using the initial set of documents, this procedure can be catastrophic when topic drift occurs. For doc2vec, one way to obtain a representation for a previously-unseen document is to keep all parameters fixed and train the representation of the new document using the standard training algorithm [208]. This approach can also be used in the case of LSE or NVSM. However, there are no guarantees that the obtained representation will be of desirable quality. In addition, the same problem remains as with the bag-of-words methods. That is, the previously-mentioned incremental updating mechanism is likely to fail when topic drift occurs.

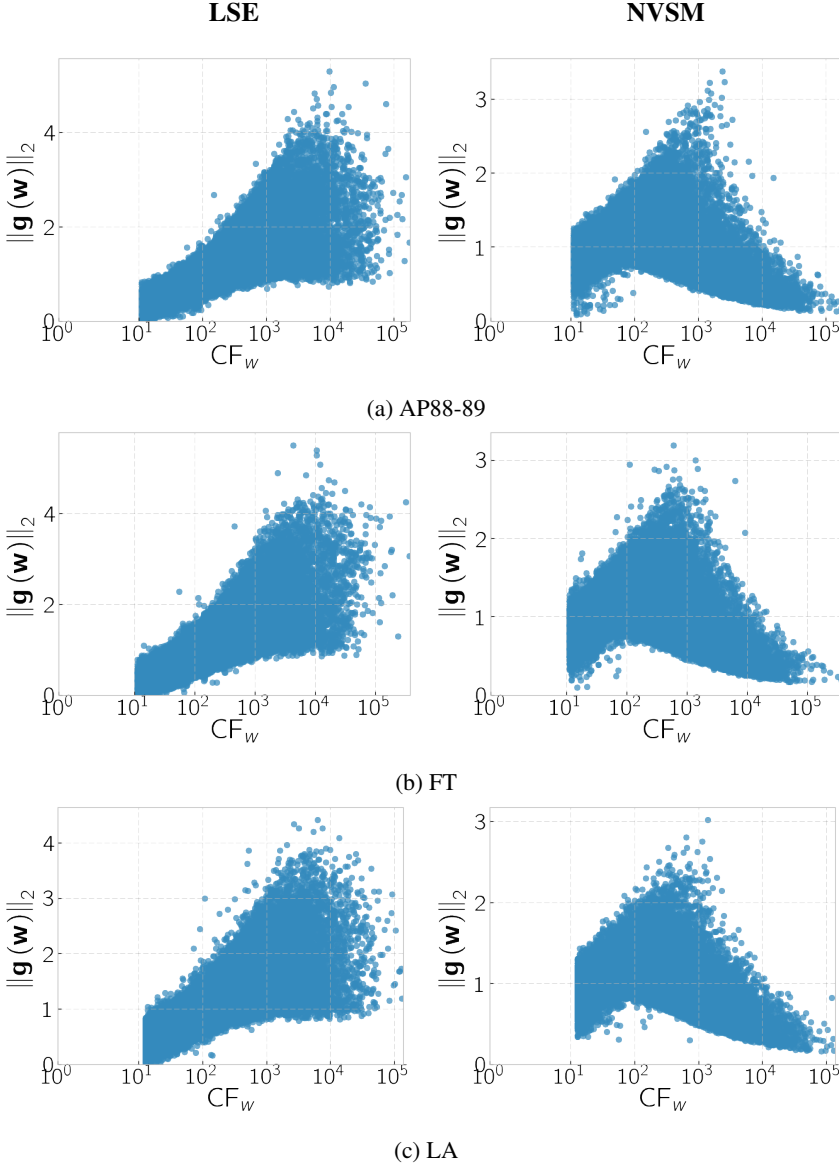


Figure 8.2: Scatter plots of term frequency in the document collections and the L2-norm of the LSE (left) and NVSM (right) representations of these terms. In the case of LSE (left scatter plots for every benchmark), we observe that the L2-norm of term representations grows linearly with the term collection frequency, and consequently, high-frequency terms are of greater importance within LSE representations. For NVSM (right scatter plot for every benchmark), we observe that terms of mid-frequency (middle 50%) have a statistically significant ($p < 0.01$) higher L2-norm (Section 8.4.4) and consequently are of greater importance for retrieval.

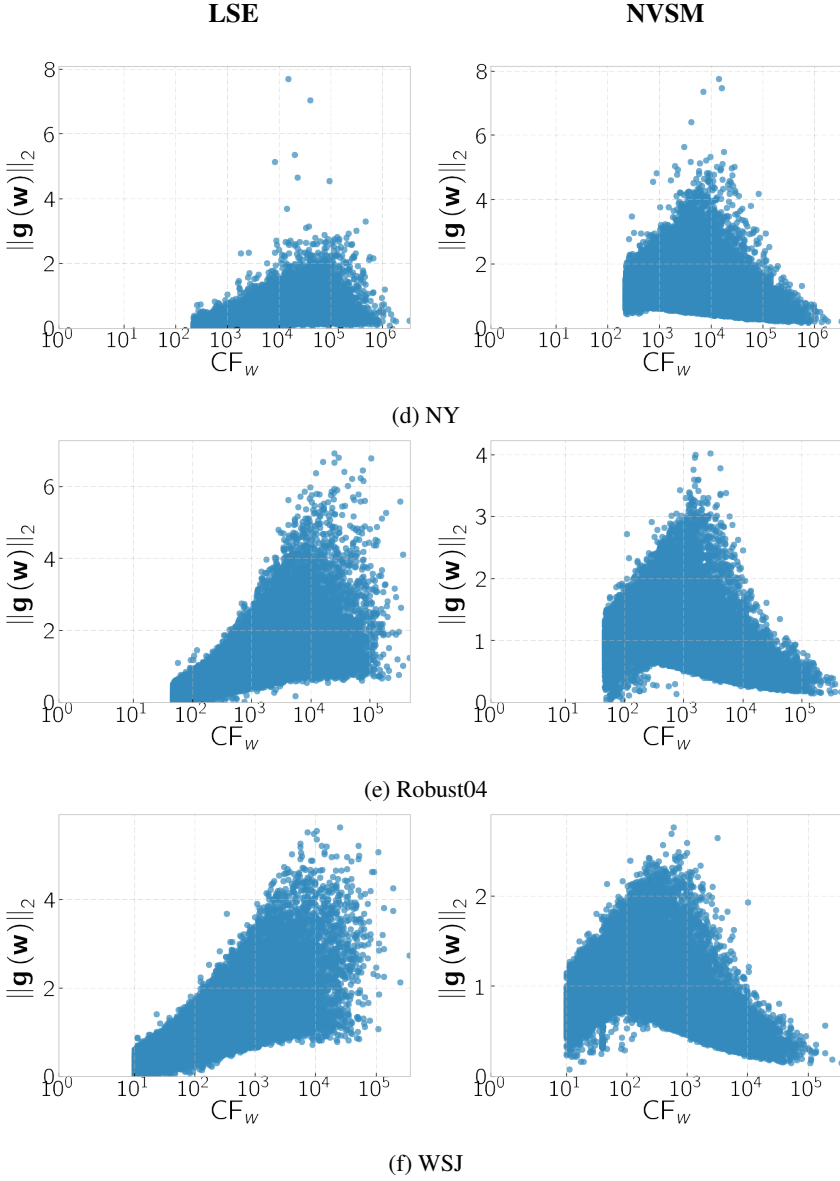


Figure 8.2 (cont'd): Scatter plots of term frequency in the document collections and the L2-norm of the LSE (left) and NVSM (right) representations of these terms. In the case of LSE (left scatter plots for every benchmark), we observe that the L2-norm of term representations grows linearly with the term collection frequency, and consequently, high-frequency terms are of greater importance within LSE representations. For NVSM (right scatter plot for every benchmark), we observe that terms of mid-frequency (middle 50%) have a statistically significant ($p < 0.01$) higher L2-norm (Section 8.4.4) and consequently are of greater importance for retrieval.

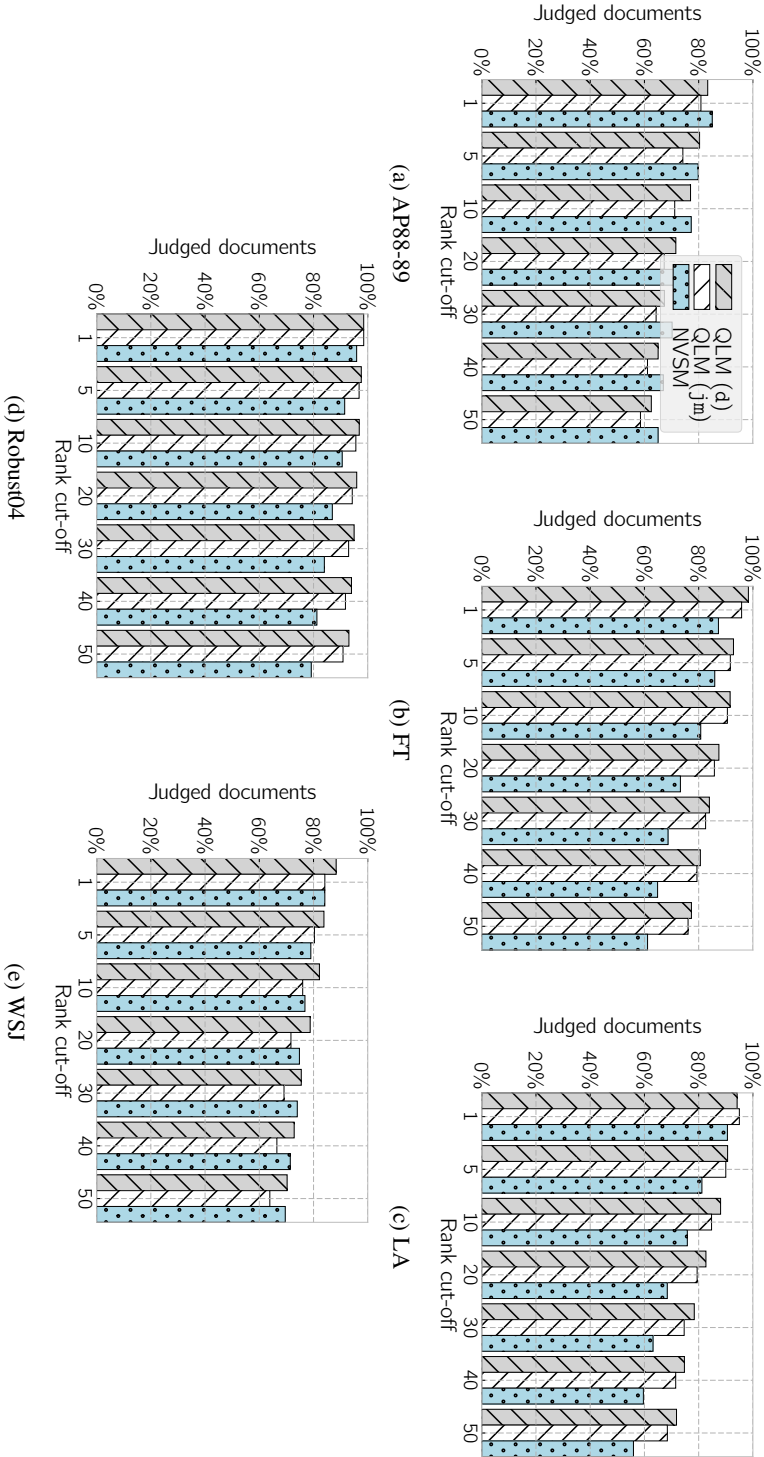


Figure 8.3: Investigation of the average number of judged document (relevant and non-relevant) at different ranks for bag-of-words methods (QLM) and NVSM on article search benchmarks (Section 8.4.2).

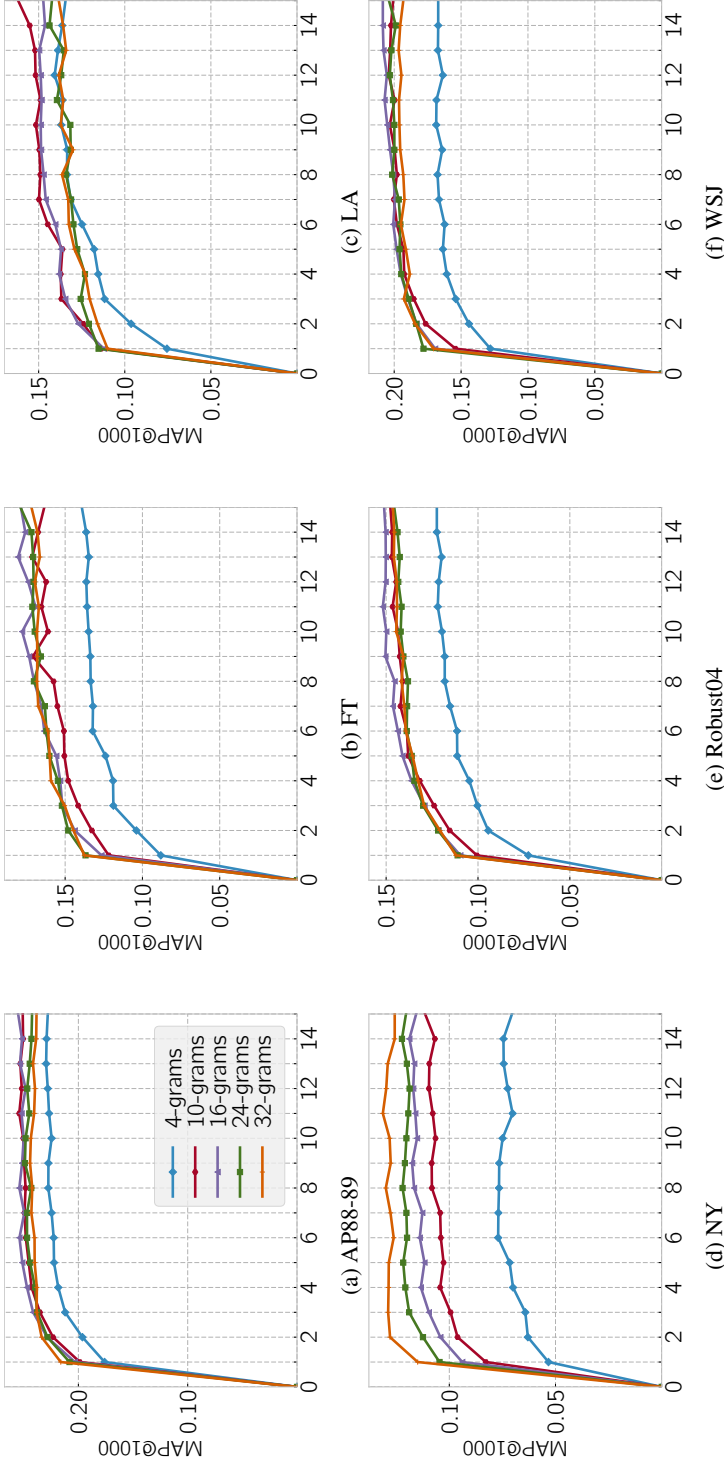


Figure 8.4: Test set MAP@1000 as training progresses on article search benchmarks with document space dimensionality $k_d = 256$. We see that MAP@1000 converges to a fixed performance level with differently-sized n -grams (here we show $n = 4, 10, 16, 24, 32$; the curves for the remaining values $n = 6, 8, 12$ are qualitatively similar and omitted to avoid clutter).

9

Conclusions

In this dissertation, we have devoted six research chapters to address two limitations of the inverted index that contribute to the vocabulary gap between query and document. This vocabulary gap causes decreased retrieval effectiveness that occurs due to: (1) the use of complex textual structures as an unfiltered query with many—possibly misleading—terms and (2) queries and their relevant documents that use different words to describe the same concepts. Specifically, Part I considers the task of formulating an effective query from a complex textual structure (i.e., search sessions, email threads) such that the formulated query is more focused and better satisfies the information need. Part II is dedicated to latent vector spaces that allow us to bridge the semantic vocabulary gap.

In this final chapter, we revisit the research questions we answered and summarize our findings in Section 9.1. In Section 9.2, we discuss limitations of our work and directions for future work.

9.1 Main findings

We now revisit our research questions introduced in Chapter 1 and summarize our findings.

RQ1 How to formulate a query from complex textual structures—such as search sessions or email threads—in order to better answer an information need?

To answer our first question, we performed an analysis of TREC search session logs in Chapter 3 and introduced a frequency-based query term weighting method that summarizes the user’s information need. We found that the semantic vocabulary gap, the subject of the next question, is prevalent in session search and that methods restricted to the user-specified query terms face a very strict performance ceiling. The focus of Chapter 4 was to formulate a query from an email thread. In particular, we focused on the case where an incoming email contains a request for content and a query needs to be formulated to retrieve the relevant attachment from a repository. The information need in the email scenario is less clear than in session search, as email messages have multiple aspects and the request for content can be implicit. We introduced a methodology for constructing a pseudo test/training collection from email collections, including

the construction of silver-standard training queries, and proposed a neural network architecture that learns to select query terms from the incoming request message.

The principal shortcoming of the query formulation methods explored in **RQ1** is that the extracted query terms are limited to the terms present in the complex textual structures. In session search, query reformulations performed by the user can succinctly describe the information need as the user directly interacts with the search engine. However, for the case of email attachment recommendation where the interaction between requester and search engine is indirect and possibly even unknown to the requester, the conditions for a serious discrepancy between the content request and the document repository have been met. Consequently, the semantic mismatch between query/document was addressed in the following questions.

RQ2 Can we learn a latent vector space of retrievable entities that performs well for retrieval?

We introduced a latent semantic model for the expert finding task in Chapter 5. We showed that our latent model, that consists of word/expert representations and a prior over experts, outperforms state-of-the-art retrieval models and contributes a complementary signal to lexical models. In Chapter 6, we investigated the structural regularities that are present in entity vector spaces. We showed that entity (i.e., expert individuals) representations estimated from associated texts alone can be used as feature vectors for clustering and recommendation. In addition, we showed that the prior over entities of the model introduced in Chapter 5 encodes entity salience.

However, the expert finding task is characterized by a few thousand entities that are each represented by a sizeable collection of documents that cover their expertise. This setting is ideal for latent semantic models like ours that learn from context due to the abundance of textual content. In addition, queries used to search in expert finding are informational [44], cover a broad topic and consequently have a need for semantic matching. Note that the case when the user knows the name of the expert (i.e., known-item search) is different from the expert finding setting. This brings us to our following question: can we adapt our latent model to larger entity domains where textual content is scarce? We answered this question as part of **RQ3**.

RQ3 Can we scale up latent vector spaces to larger entity domains that have less textual content per entity compared to the expert finding setting?

Chapter 7 introduced Latent Semantic Entities (LSE), a modification of our latent model for expert finding where queries and entities are represented in a latent metric space. This implies that we no longer learn an importance prior over entities. In addition, the scalability of the learning mechanism was improved by sampling. We evaluated LSE in a product search setting and showed that it improves retrieval effectiveness as a complementary signal next to product saliency and lexical matching in a product-oriented search engine.

The solution to **RQ3** was a training mechanism that is based on sampling rather than considering the full set of entities. Nearest neighbour algorithms can be used to rank entities with a time complexity sub-linear w.r.t. the number of entities. However,

entity ranking is only a small part of information retrieval. How do our latent vector spaces operate in a more traditional setting, such as news article retrieval?

RQ4 Can we further increase the retrieval effectiveness and scale latent vector spaces up to hundreds of thousands of documents?

With Chapter 8, we saw the introduction of an extension to LSE: the Neural Vector Space Model (NVSM). We evaluated NVSM on article retrieval benchmarks from TREC and showed that it outperforms all of the existing state-of-the-art latent vector spaces. In addition, we showed that NVSM significantly improves retrieval effectiveness when added as a complimentary feature in addition to a lexical language model and another latent vector space model. Consequently, NVSM contributes an additional signal. A comparative analysis showed that NVSM performs better on queries than other latent vector spaces when the queries require a greater extent of lexical matching. Further investigation resulted in the observation that NVSM learns a notion of Luhn significance, a quantity known to be important for retrieval. In particular, we found that the L2-norm of mid-frequency words is significantly larger (and thus, the words are of greater importance) than low- and high-frequency words.

To conclude this section, we reflect on the first chapter of this dissertation and repeat the two major drawbacks of the inverted index that motivated our research: (1) Queries consisting of many terms induce high computational costs, while there often exists a shorter query that is more effective in fulfilling an information need. We addressed this drawback in Part I of this dissertation by answering **RQ1**. In particular, we focused on the case where we wish to formulate a query from a complex textual structure. However, we also found that term-based matching by itself is not sufficient to fulfil information needs as there exists a vocabulary gap between the user query and the relevant document. (2) Part II of this dissertation was dedicated to bridging the vocabulary gap. The use of term-based matching to build an initial candidate set of documents may incorrectly classify relevant documents that do not contain query terms as irrelevant. We addressed this issue in **RQ2**, **RQ3** and **RQ4** by the development of latent vector spaces where queries and documents are matched according to their semantics rather than exact term occurrences. Consequently, term-based retrieval may be complemented by semantic vector spaces that can be queried through a nearest neighbour search in a low-dimensional vector space.

9.2 Future work

This dissertation resulted in insights and algorithms for bridging the vocabulary gap in IR. However, the research performed as part of this dissertation raised more questions than it answered. In this section, we summarize the limitations of our work and conclude this dissertation with directions for future work.

The limitations of Part I of this dissertation are as follows. (a) In Part I we only considered terms occurring within the complex textual structures as candidates. While this prevents the term candidate set from becoming too large, it does limit the ability for methods to formulate expressive queries in the case where textual data is scarce. (b) The

retrieval model used in Part I, a language model with Dirichlet smoothing, is ubiquitous in retrieval systems. However, smoothing allows the search engine to deal with verbose queries [286] that contain terms absent from the messages. Consequently, our findings may change when considering other retrieval model classes, such as boolean models or semantic matching models.

Considering Part II of this dissertation, its limitations are as follows: (a) The largest retrieval collection used in Part II of this dissertation consists of half a million documents. While we obtained promising results on larger collections of up to two million documents (not included in this dissertation), the question remains of whether latent vector spaces are applicable in large retrieval scenarios. The question of applicability applies to two separate aspects. The first aspect pertains to the training of latent vector space models. More specifically, the training of latent vector space models is limited by their space complexity that grows linearly with the number of document terms. Secondly, the question remains whether (approximate) nearest neighbour search algorithms are actually efficient enough to perform retrieval in real-time. This brings us to the second limitation of Part II. (b) In this dissertation, we assumed that the modelling of latent vector spaces is separated from the development of nearest neighbour algorithms that are used to query them. Particularly, we did not evaluate the effect of using approximate nearest neighbour algorithms—which are likely required for real-time querying—on retrieval effectiveness. However, in this dissertation we performed our retrieval evaluation using offline test collections. In addition, the focus of this dissertation lies on the modelling side of things. Consequently, the use of exact nearest neighbour algorithms is justified in this dissertation.

To address these limitations, we identify the following directions for future work:

Query formulation. We explored the task of formulating queries from complex textual structures, with applications to session search and email. We first discuss directions for the applications and then focus on the general task. In the case of session search (Chapter 3), there is still much room for improvement by re-weighting query terms. Future work should focus on better lexical query models for session search, in addition to semantic matching and tracking the dynamics of contextualized semantics in search.

For email attachment recommendation (Chapter 4), future work includes the incorporation of social connections in the email domain where a social graph can be constructed from email interactions and entity mentions. In addition, structured queries with operators searching different fields (e.g., recipients, subject) can improve performance. Finally, we assumed a single model for all mailboxes. However, per-mailbox specialized models are likely to generate better queries. Overall, the query/document mismatch is prevalent when formulating queries from complex textual structures and methods restricted to lexical query modelling face a very strict performance ceiling. Consequently, future work should focus on formulating queries using the full set of terms, instead of only those occurring in the structures. This is non-trivial as retrieval systems deal with very large dictionaries, and therefore, the effect of including a single term is hard to estimate.

Latent representations as feature vectors. In Chapter 6 we explored using latent entity representations as feature vectors in other applications, such as clustering, recommendation and determining entity salience. Future work includes the use of text-based entity representations in end-to-end applications. For example, in social networks these methods can be applied to cluster users or to induce graphs based on thread participation or hashtag usage. In addition, text-based entity representations can be used as item feature vectors in recommendation systems. Beyond text-only entity collections, there is also a plenitude of applications where entity relations are available. While there has been some work on learning latent representations from entity relations [39, 289], little attention has so far been given to combining textual evidence and entity relations. Therefore, we identify two additional directions for future work. First, an analysis showing in what capacity entity representations estimated from text alone encode entity-entity relations (beyond the co-associations considered in this work). Secondly, the incorporation of entity-entity similarity in the construction of latent entity representations.

Latent vector spaces for information retrieval. Part II of this dissertation is centered around the construction of latent vector spaces for the retrieval task. While significant progress was made, more questions and directions for future work arise: (1) How can we further improve the retrieval effectiveness of latent vector spaces? Future work includes adding additional model expressiveness through depth or width. In addition, multiple nearest neighbour searches for every query term could be performed in parallel and consequently, the influence of individual query terms can be combined in a more expressive way. (2) Additional signals of relevance (e.g., query/document pairs) or similarity, such as entity/entity similarity [209], can be incorporated during training. Modelling objects beyond entities and query terms, such as users within a personalization context [6] is also a promising direction. (3) Can we scale up the latent vector spaces to hundreds of millions of documents? This direction introduces both engineering and modelling challenges. One way to scale up the existing vector spaces presented in this dissertation is to train multiple models in parallel on sub-samples of the full document collection. (4) How do our latent vector spaces perform in online settings? In particular, what is the effect of approximate nearest neighbour algorithms on retrieval performance? This direction comes with non-trivial engineering challenges and is best performed when one has access to a platform with actual users.

Appendices



Pyndri: A Python Interface to the Indri Search Engine

A.1 Introduction

Research in Artificial Intelligence progresses at a rate proportional to the time it takes to implement an idea. Therefore, it is natural for researchers to prefer scripting languages (e.g., Python) over conventional programming languages (e.g., C++) as programs implemented using the latter are often up to three factors longer (in lines of code) and require twice as much time to implement [204]. Python, an interactive scripting language that emphasizes readability, has risen in popularity due to its wide range of scientific libraries (e.g., NumPy), built-in data structures and holistic language design [135].

There is still, however, a lack of an integrated Python library dedicated to information retrieval research. Researchers often implement their own procedures to parse common file formats, perform tokenization, token normalization that encompass the overall task of corpus indexing. Uysal and Gunal [249] show that text classification algorithms can perform significantly differently, depending on the level of preprocessing performed. Existing frameworks, such as NLTK [153], are primarily targeted at processing natural language as opposed to retrieving information and do not scale well. At the algorithm level, small implementation differences can have significant differences in retrieval performance due to floating point errors [92]. While this is unavoidable due to the fast-paced nature of research, at least for seminal algorithms and models, standardized implementations are needed.

A.2 Introducing Pyndri

Fortunately, the IR community has developed a series of indexing frameworks (e.g., Galago, Lucene, Terrier) that correctly implement a wide range of retrieval models. The Indri search engine [238] supports complex queries involving *evidence combination* and the ability to specify a wide variety of constraints involving *proximity*, *syntax*, *extracted entities* and *document structure*. Furthermore, the framework has been efficiently implemented using C++ and was designed from the ground up to support *very large*

```
index = pyndri.Index('/opt/local/clueweb09')

for int_doc_id in range(index.document_base(),
                        index.maximum_document()):
    ext_doc_id, doc_tokens = index.document(int_doc_id)
```

Code snippet 1: Tokenized documents in the index can be iterated over. The `ext_doc_id` variable in the inner loop will equal the document identifier (e.g., `clueweb09-en0039-05-00000`), while the `doc_tokens` points to a tuple of integers that correspond to the document term identifiers.

databases, optimized query execution and fast and concurrent indexing. A large subset of the retrieval models [20, 29, 100, 258, 285] introduced over the course of history can be succinctly formulated as an Indri query. However, to do so in an automated manner, up until now researchers were required to resort to C++, Java or shell scripting. C++ and Java, while excellent for production-style systems, are slow and inflexible for the fast prototyping paradigm used in research. Shell scripting fits better in the research paradigm, but offers poor string processing functionality and can be error-prone. Besides, shell scripting is unsuited if one wants to evaluate a large number of complex queries or wishes to extract documents from the repository as this incurs overhead, causing avoidable slow execution. Existing Python libraries for indexing and searching, such as PyLucene, Whoosh or Elasticsearch, do not support the rich Indri language and functionality required for rapid prototyping.

We fill this gap by introducing `pyndri`, a lightweight interface to the Indri search engine. `Pyndri` offers read-only access at two levels in a given Indri index.

A.2.1 Low-level access to document repository

First of all, `pyndri` allows the retrieval of tokenized documents stored in the index repository. This allows researchers to avoid implementing their own format parsing as Indri supports all major formats used in IR, such as the `trectext`, `trecweb`, XML documents and Web ARChive (WARC) formats. Furthermore, standardized tokenization and normalization of texts is performed by Indri and is no longer a burden to the researcher. Code snippet 1 shows how a researcher can easily access documents in the index. Lookup of internal document identifiers given their external name is provided by the `Index.document_ids` function.

The dictionary of the index (Code snippet 2) can be accessed from Python as well. Beyond bi-directional token-to-identifier translation, the dictionary contains corpus statistics such as term and document frequencies as well. The combination of index iteration and dictionary interfacing integrates conveniently with the Gensim¹ package, a collection of topic and latent semantic models such as LSI [70] and word2vec [177]. In particular for word2vec, this allows for the training of word embeddings on a corpus while avoiding the tokenization mismatch between the index and word2vec.

¹<https://radimrehurek.com/gensim>

```
index = pyndri.Index('/opt/local/clueweb09')
dictionary = pyndri.extract_dictionary(index)

_, int_doc_id = index.document_ids(
    ['clueweb09-en0039-05-00000'])
print([dictionary[token_id]
       for token_id in index.document(int_doc_id)[1]])
```

Code snippet 2: A specific document is retrieved by its external document identifier. The index dictionary can be queried as well. In the above example, a list of token strings corresponding to the document's contents will be printed to `stdout`.

```
index = pyndri.Index('/opt/local/clueweb09')

for int_doc, score in index.query('obama family tree'):
    # Do stuff with the document.
```

Code snippet 3: Simple queries can be fired using a simple interface. Here we query the index for topic `wt09-1` from the TREC 2009 Web Track using the Indri defaults (Query Language Model (QLM) with Dirichlet smoothing, $\mu = 2500$).

In addition to tokenized documents, `pyndri` also supports retrieving various corpus statistics such as document length and corpus term frequency.

A.2.2 Querying Indri from Python

Secondly, `pyndri` allows the execution of Indri queries using the index. Code snippet 3 shows how one would query an index using a topic from the TREC 2009 Web Track using the Indri default retrieval model. Beyond simple terms, the `query()` function fully supports the Indri Query Language.²

In addition, we can specify a subset of documents to query, the number of requested results and whether or not snippets should be returned. In Code snippet 4 we create a `QueryEnvironment`, with a set of custom smoothing rules. This allows the user to apply fine-grained smoothing settings (i.e., per-field granularity).

A.3 Summary

In this appendix we introduced `pyndri`, a Python interface to the Indri search engine. `Pyndri` allows researchers to access tokenized documents from Indri using a convenient Python interface. By relying on Indri for tokenization and normalization, IR researchers are no longer burdened by this task. In addition, complex retrieval models can easily be

²<http://lemurproject.org/lemur/IndriQueryLanguage.php>

```
index = pyndri.Index('/opt/local/clueweb09')
query_env = pyndri.QueryEnvironment(
    index, rules=('method:dirichlet,mu:5000',))

results = query_env.query(
    '#weight( 0.70 obama 0.20 family 0.10 tree )',
    document_set=map(
        operator.itemgetter(1),
        index.document_ids([
            'clueweb09-en0003-55-31884',
            'clueweb09-en0006-21-20387',
            'clueweb09-enwp01-75-20596',
            'clueweb09-enwp00-64-03709',
            'clueweb09-en0005-76-03988'
        ])),
    results_requested=3,
    include_snippets=True)

for int_doc_id, score, snippet in results:
    # Do stuff with the document and snippet.
```

Code snippet 4: Advanced querying of topic wt09-1 with custom smoothing rules, using a weighted-QLM. Only a subset of documents is searched and we impose a limit on the size of the returned list. In addition to the document identifiers and their retrieval score, the function now returns snippets of the documents where the query terms match.

implemented by constructing them in the Indri Query Language in Python and querying the index. This will make it easier for researchers to release their code, as Python is designed to be readable and cross-platform. We hope that with the release of pyndri, we will stimulate **reproducible**, **open** and **fast-paced** IR research. More information regarding the available API and installation instructions can be found on Github.³

³<https://github.com/cvangysel/pyndri>

B

Semantic Entity Retrieval Toolkit

B.1 Introduction

The unsupervised learning of low-dimensional, semantic representations of words and entities has recently gained attention for the entity-oriented tasks of expert finding [257] and product search [256]. Representations are learned from a document collection and domain-specific associations between documents and entities. Expert finding is the task of finding the right person with the appropriate skills or knowledge [23] and an association indicates document authorship (e.g., academic papers) or involvement in a project (e.g., annual progress reports). In the case of product search, an associated document is a product description or review [256].

In this appendix we describe the Semantic Entity Retrieval Toolkit (SERT) that provides implementations of our previously published entity representation models [256, 257]. Beyond a unified interface that combines different models, the toolkit allows for fine-grained parsing configuration and GPU-based training through integration with Theano [73, 242]. Users can easily extend existing models or implement their own models within the unified framework. After model training, SERT can compute matching scores between an entity and a piece of text (e.g., a query). This matching score can then be used for ranking entities, or as a feature in a downstream machine learning system, such as the learning to rank component of a search engine. In addition, the learned representations can be extracted and used as feature vectors in entity clustering or recommendation tasks [259]. The toolkit is licensed under the permissive MIT open-source license.¹

B.2 The toolkit

SERT is organized as a pipeline of utilities as depicted in Fig. B.1. First, a collection of documents and entity associations is processed and packaged using a numerical format (Section B.2.1). Low-dimensional representations of words and entities are then learned (Section B.2.2) and afterwards the representations can be used to make inferences (Section B.2.3).

¹The toolkit is licensed under the permissive MIT open-source license and can be found at <https://github.com/cvangysel/SERT>.




Figure B.1: Schematic overview of the different pipeline components of SERT. The collection is parsed, processed and packaged in a numerical format using the *prepare* (Section B.2.1) utility. Afterwards, the *training* (Section B.2.2) utility learns representations of entities and words and the *query* (Section B.2.3) utility is used to compute matching scores between entities and queries.

B.2.1 Collection parsing and preparation

To begin, SERT constructs a vocabulary that will be used to tokenize the document collection. Non-significant words that are too frequent (e.g., stopwords), noisy (e.g., single characters) and rare words are filtered out. Words that do not occur in the dictionary are ignored. Afterwards, word sequences are extracted from the documents and stored together with the associated entities in the numerical format provided by NumPy [250]. Word sequences can be extracted consecutively or a stride can be specified to extract non-consecutive windows. In addition, a hierarchy of word sequence extractors can be applied to extract skip-grams, i.e., word sequences where a number of tokens are skipped after selecting a token [105]. To support short documents, a special-purpose padding token can be used to fill up word sequences that are longer than a particular document.

After word sequence extraction, a weight can be assigned to each word sequence/entity pair that can be used to re-weight the training objective. For example, in the case of expert finding [257], this weight is the reciprocal of the document length of the document where the sequence was extracted from. This avoids a bias in the objective towards long documents. An alternative option that exists within the toolkit is to resample word sequence/entity pairs such that every entity is associated with the same number of word sequences, as used for product search [256].

B.2.2 Representation learning

After the collection has been processed and packaged in a machine-friendly format, representations of words and entities can be learned. The toolkit includes implementations of state-of-the-art representation learning models that were applied to expert finding [257] and product search [256]. Users of the toolkit can use these implementations to learn representations out-of-the-box or adapt the algorithms to their needs. In addition, users can implement their own models by extending an interface provided by the framework. Code snippet 5 shows an example of a model implemented in the SERT toolkit where users can define a symbolic cost function that will be optimized using Theano [242]. Due to the component-wise organization of the toolkit (Fig. B.1), modelling and text processing are separated from each other. Consequently, researchers can focus on modelling and representation learning only. In addition, any improvements

```

class ExampleModel(VectorSpaceLanguageModelBase):

    def __init__(self, *args, **kwargs):
        super(ExampleModel, self).__init__(
            *args, **kwargs)

        # Define model architecture.
        input_layer = InputLayer(
            shape=(self.batch_size, self.window_size))

        ...

    def loss_fn(pred, actual, _):
        # Compute symbolic loss between
        # predicted/actual entities.

        # The framework deals with underlying boilerplate.
        self._finalize(loss_fn, ...)

    def get_representations(self):
        # Returns the representations and parameters
        # to be extracted.

```

Code snippet 5: Illustrative example of the SERT model interface. The full interface supports more functionality omitted here for brevity. Users can define a symbolic graph of computation using the Theano library [242] in combination with Lasagne [73].

to the collection processing (Section B.2.1) collectively benefits all models implemented in SERT.

B.2.3 Entity ranking & other uses of the representations

Once a model has been trained, SERT can be used to rank entities w.r.t. a textual query. The concrete implementation used to rank entities depends on the model that was trained. In the most generic case, a matching score is computed for every entity and entities are ranked in decreasing order of his score. However, in the special case when the model is interpreted as a metric vector space [43, 256], SERT casts entity ranking as a k -nearest neighbour problem and uses specialized data structures for retrieval [130]. After ranking, SERT outputs the entity rankings as a TREC-compatible file that can be used as input to the `trec_eval`² evaluation utility.

Apart from entity ranking, the learned representations and model-specific parameters can be extracted conveniently from the models through the interface³ and used for downstream tasks such as clustering, recommendation and determining entity importance as

²https://github.com/usnistgov/trec_eval

³See `get_representations` in Snippet 5.

shown in [259].

B.3 Summary

In this appendix we described the Semantic Entity Retrieval Toolkit, a toolkit that learns latent representations of words and entities. The toolkit contains implementations of state-of-the-art entity representations algorithms [256, 257] and consists of three components: text processing, representation learning and inference. Users of the toolkit can easily make changes to existing model implementations or contribute their own models by extending an interface provided by the SERT framework.

Future work includes integration with Pyndri [262] such that document collections indexed with Indri can transparently be used to train entity representations. In addition, integration with machine learning frameworks besides Theano, such as TensorFlow and PyTorch, will make it easier to integrate existing models into SERT.

Bibliography

- [1] The knowledge-based economy. Technical report, Organisation for Economic Co-operation and Development, 1996. (Cited on page 51.)
- [2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. (Cited on page 114.)
- [3] Q. Ai, L. Yang, J. Guo, and W. B. Croft. Analysis of the paragraph vector model for information retrieval. In *ICTIR*, pages 133–142. ACM, 2016. (Cited on page 13.)
- [4] Q. Ai, L. Yang, J. Guo, and W. B. Croft. Improving language estimation with the paragraph vector model for ad-hoc retrieval. In *SIGIR*, pages 869–872. ACM, 2016. (Cited on pages 13, 108, and 116.)
- [5] Q. Ai, S. T. Dumais, N. Craswell, and D. Liebling. Characterizing email search using large-scale behavioral logs and surveys. In *WWW*, pages 1511–1520, 2017. (Cited on pages 10 and 29.)
- [6] Q. Ai, Y. Zhang, K. Bi, X. Chen, and B. W. Croft. Learning a hierarchical embedding model for personalized product search. In *SIGIR*, 2017. (Cited on page 139.)
- [7] J. Allan, B. Croft, A. Moffat, and M. Sanderson. Frontiers, challenges, and opportunities for information retrieval. In *SIGIR Forum*, volume 46, pages 2–32. ACM, 2012. (Cited on page 10.)
- [8] J. Allan, D. Harman, E. Kanoulas, D. Li, C. Van Gysel, and E. Vorhees. Trec 2017 common core track overview. In *TREC*, 2017. (Cited on pages 115 and 116.)
- [9] E. Amigó, J. Gonzalo, J. Artiles, and F. Verdejo. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval*, 12(4):461–486, 2009. ISSN 1386-4564. (Cited on page 80.)
- [10] J. Arguello, S. Avula, and F. Diaz. Using query performance predictors to improve spoken queries. In *ECIR*, pages 309–321. Springer, 2016. (Cited on page 12.)
- [11] N. Asadi, D. Metzler, T. Elsayed, and J. Lin. Pseudo test collections for learning web search ranking functions. In *SIGIR*, pages 1073–1082. ACM, 2011. (Cited on page 34.)
- [12] L. Azzopardi, M. de Rijke, and K. Balog. Building simulated queries for known-item topics: An analysis using six european languages. In *SIGIR*. ACM, 2007. (Cited on page 34.)
- [13] P. Bailey, A. P. de Vries, N. Craswell, and I. Soboroff. Overview of the TREC 2007 enterprise track. In *TREC*, 2007. (Cited on page 56.)
- [14] N. Balasubramanian, G. Kumaran, and V. R. Carvalho. Exploring reductions for long web queries. In *SIGIR*, pages 571–578. ACM, 2010. (Cited on page 11.)
- [15] K. Balog. *People Search in the Enterprise*. PhD thesis, University of Amsterdam, 2008. (Cited on page 56.)
- [16] K. Balog. On the investigation of similarity measures for product resolution. In *LHD workshop at IJCAI*, 2011. (Cited on page 14.)
- [17] K. Balog and M. de Rijke. Determining expert profiles (with an application to expert finding). In *IJCAI*, 2007. (Cited on page 88.)
- [18] K. Balog and M. de Rijke. Finding similar experts. In *SIGIR*, pages 821–822. ACM, 2007. (Cited on pages 14 and 82.)
- [19] K. Balog and R. Neumayer. A test collection for entity search in dbpedia. In *SIGIR*, pages 737–740. ACM, 2013. (Cited on page 87.)
- [20] K. Balog, L. Azzopardi, and M. de Rijke. Formal models for expert finding in enterprise corpora. In *SIGIR*, pages 43–50. ACM, 2006. (Cited on pages 2, 14, 33, 42, 52, 56, 58, 64, 65, 66, 69, 77, 82, 84, 85, 98, and 144.)
- [21] K. Balog, L. Azzopardi, and M. de Rijke. A language modeling framework for expert finding. *IPM*, 45:1–19, 2009. (Cited on pages 14, 51, and 58.)
- [22] K. Balog, P. Serdyukov, and A. P. de Vries. Overview of the TREC 2010 entity track. Technical report, DTIC Document, 2011. (Cited on pages 13 and 87.)
- [23] K. Balog, Y. Fang, M. de Rijke, P. Serdyukov, and L. Si. Expertise retrieval. *Found. & Tr. in Information Retrieval*, 6(2-3):127–256, 2012. (Cited on pages 3, 13, 14, 51, 52, 56, 74, 87, and 147.)
- [24] M. Baroni, G. Dinu, and G. Kruszewski. Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL*, pages 238–247, 2014. (Cited on pages 15, 73, 88, and 107.)
- [25] I. Becerra-Fernandez. Role of artificial intelligence technologies in the implementation of People-

B. Bibliography

- Finder knowledge management systems. *Knowledge-Based Systems*, 13(5):315–320, 2000. (Cited on pages 13 and 51.)
- [26] S. M. Beitzel, E. C. Jensen, A. Chowdhury, and D. Grossman. Using titles and category names from editor-driven taxonomies for automatic evaluation. In *CIKM*, pages 17–23. ACM, 2003. (Cited on page 34.)
- [27] M. Bendersky and W. B. Croft. Discovering key concepts in verbose queries. In *SIGIR*, pages 491–498. ACM, 2008. (Cited on pages 11, 37, and 39.)
- [28] M. Bendersky and W. B. Croft. Analysis of long queries in a large scale search log. In *Workshop on Web Search Click Data*, pages 8–14. ACM, 2009. (Cited on page 32.)
- [29] M. Bendersky, D. Metzler, and W. B. Croft. Learning concept importance using a weighted dependence model. In *WSDM*, pages 31–40. ACM, 2010. (Cited on page 144.)
- [30] M. Bendersky, D. Metzler, and W. B. Croft. Effective query formulation with multiple information sources. In *SIGIR*, pages 443–452. ACM, 2012. (Cited on page 23.)
- [31] J. R. Benetka, K. Balog, and K. Nørkvåg. Anticipating information needs based on check-in activity. In *WSDM*, pages 41–50. ACM, 2017. (Cited on pages 10 and 29.)
- [32] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin. A neural probabilistic language model. *JMLR*, 3: 1137–1155, 2003. (Cited on pages 15, 54, 73, and 107.)
- [33] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. (Cited on page 15.)
- [34] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *JSTOR*, pages 289–300, 1995. (Cited on pages 59 and 67.)
- [35] P. N. Bennett, R. W. White, W. Chu, S. T. Dumais, P. Bailey, F. Borisyuk, and X. Cui. Modeling the impact of short- and long-term behavior on search personalization. In *SIGIR*, pages 185–194. ACM, 2012. (Cited on page 25.)
- [36] R. Berendsen, M. de Rijke, K. Balog, T. Bogers, and A. van den Bosch. On the assessment of expertise profiles. *JASIST*, 64(10):2024–2044, 2013. (Cited on pages 51, 56, 77, and 84.)
- [37] R. Berendsen, M. Tsagkias, W. Weerkamp, and M. de Rijke. Pseudo test collections for training and tuning microblog rankers. In *SIGIR*, pages 53–62. ACM, 2013. (Cited on page 34.)
- [38] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *JMLR*, 3:993–1022, 2003. (Cited on pages 12, 75, 91, 97, 107, 108, 114, and 118.)
- [39] A. Bordes, J. Weston, R. Collobert, and Y. Bengio. Learning structured embeddings of knowledge bases. In *AAAI*, 2011. (Cited on pages 15, 73, 88, and 139.)
- [40] A. Borisov, I. Markov, M. de Rijke, and P. Serdyukov. A neural click model for web search. In *WWW*, pages 531–541. International World Wide Web Conferences Steering Committee, 2016. (Cited on page 16.)
- [41] A. Borisov, I. Markov, M. de Rijke, and P. Serdyukov. A context-aware time model for web search. In *SIGIR*, pages 205–214. ACM, 2016. (Cited on page 16.)
- [42] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*, pages 177–186. Springer, 2010. (Cited on page 69.)
- [43] L. Boytsov, D. Novak, Y. Malkov, and N. Eric. Off the beaten path: Let’s replace term-based retrieval with k-nn search. In *CIKM*, 2016. (Cited on pages 2, 107, 116, and 149.)
- [44] A. Broder. A taxonomy of web search. *SIGIR forum*, 36(2):3–10, 2002. (Cited on pages 2, 120, and 136.)
- [45] C. Buckley, D. Dimmick, I. Soboroff, and E. Voorhees. Bias and the limits of pooling for large collections. *Information retrieval*, 10(6):491–508, 2007. (Cited on pages 119 and 125.)
- [46] J. Budzik and K. Hammond. Watson: Anticipating and contextualizing information needs. In *ASIS*, volume 36, pages 727–740. Information Today, 1999. (Cited on page 10.)
- [47] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *ICML*, pages 89–96. ACM, 2005. (Cited on page 16.)
- [48] R. Cai, H. Wang, and J. Zhang. Learning entity representation for named entity disambiguation. In *Chin. Comp. Ling. and Nat. Lang. Proc. Based on Nat. Ann. Big Data*, pages 267–278. Springer, 2015. (Cited on page 88.)
- [49] Y. Cao, J. Liu, S. Bao, and H. Li. Research on Expert Search at Enterprise Track of TREC 2005. In *TREC*, pages 2–5, 2005. (Cited on page 14.)
- [50] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In *ICML*, pages 129–136. ACM, 2007. (Cited on page 40.)
- [51] B. Carterette, E. Kanoulas, M. M. Hall, and P. D. Clough. Overview of the trec 2014 session track. In

-
- TREC*, 2014. (Cited on pages 9 and 20.)
- [52] V. R. Carvalho and W. W. Cohen. On the collective classification of email speech acts. In *SIGIR*, pages 345–352. ACM, 2005. (Cited on page 11.)
 - [53] D. D. Castro, L. Lewin-eytan, Z. Karnin, and Y. Maarek. You’ve got mail, and here is what you could do with it! In *WSDM*, 2016. (Cited on page 10.)
 - [54] S. Cetintas and L. Si. Effective query generation and postprocessing strategies for prior art patent search. *JASIST*, 63(3):512–527, 2012. (Cited on pages 11, 37, 39, and 42.)
 - [55] O. Chapelle and Y. Zhang. A dynamic bayesian network click model for web search ranking. In *WWW*, pages 1–10. ACM, 2009. (Cited on page 111.)
 - [56] M. Chen. Efficient vector representation for documents through corruption. 2017. (Cited on page 108.)
 - [57] K. Clark and C. D. Manning. Improving coreference resolution by learning entity-level distributed representations. arXiv 1606.01323, 2016. (Cited on page 73.)
 - [58] R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*, pages 160–167, 2008. (Cited on page 73.)
 - [59] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa. Natural language processing (almost) from scratch. *JMLR*, 12(Aug):2493–2537, 2011. ISSN 1532-4435. (Cited on pages 15, 16, 88, and 107.)
 - [60] G. V. Cormack, M. D. Smucker, and C. L. Clarke. Efficient and effective spam filtering and re-ranking for large web datasets. *Information retrieval*, 14(5):441–465, 2011. (Cited on page 22.)
 - [61] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012. (Cited on page 118.)
 - [62] I. B. Crabtree, S. J. Soltysiak, and M. Thint. Adaptive personal agents. *Personal Technologies*, 2(3): 141–151, 1998. (Cited on page 10.)
 - [63] N. Craswell, D. Hawking, A.-M. Vercoustre, and P. Wilkins. P@nopic expert: Searching for experts not just for documents. In *Ausweb Poster Proceedings*, pages 21–25, 2001. (Cited on pages 13 and 51.)
 - [64] N. Craswell, A. P. de Vries, and I. Soboroff. Overview of the TREC 2005 enterprise track. In *TREC*, 2005. (Cited on pages 56 and 77.)
 - [65] N. Craswell, W. B. Croft, J. Guo, B. Mitra, and M. de Rijke. Neu-ir: The sigir 2016 workshop on neural information retrieval. In *SIGIR*, pages 1245–1246. ACM, 2016. (Cited on pages 16, 107, and 109.)
 - [66] B. Croft, D. Metzler, and T. Strohman. *Search Engines: Information Retrieval in Practice*. 2015. URL <http://ciir.cs.umass.edu/downloads/SEIRiP.pdf>. (Cited on pages 1, 2, and 9.)
 - [67] R. Cummins. A study of retrieval models for long documents and queries in information retrieval. In *WWW*, pages 795–805, 2016. (Cited on page 3.)
 - [68] T. H. Davenport and L. Prusak. *Working knowledge: How organizations manage what they know*. Harvard Business Press, 1998. (Cited on page 51.)
 - [69] A. P. de Vries, A.-M. Vercoustre, J. A. Thom, N. Craswell, and M. Lalmas. Overview of the INEX 2007 entity ranking track. In *Focused Access to XML Documents*, pages 245–251. Springer, 2007. (Cited on pages 2 and 87.)
 - [70] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407, 1990. (Cited on pages 12, 75, 91, 97, 107, 108, 109, 114, 118, and 144.)
 - [71] G. Demartini, J. Gaugaz, and W. Nejdl. A vector space model for ranking entities and its application to expert search. In *ECIR*, pages 189–201. Springer, 2009. (Cited on pages 14, 51, 56, 68, 73, 75, 77, 85, 87, 88, 91, and 97.)
 - [72] L. Deng, X. He, and J. Gao. Deep stacking networks for information retrieval. In *ICASSP*, pages 3153–3157, 2013. (Cited on page 16.)
 - [73] S. Dieleman, J. Schlüter, C. Raffel, E. Olson, S. K. Sønderby, D. Nouri, et al. Lasagne: First release., Aug. 2015. URL <http://dx.doi.org/10.5281/zenodo.27878>. (Cited on pages 147 and 149.)
 - [74] L. Dietz, A. Kotov, and E. Meij. Utilizing knowledge bases in text-centric information retrieval. In *ICTIR*, pages 5–5. ACM, 2016. (Cited on page 13.)
 - [75] D. Donato, F. Bonchi, T. Chi, and Y. Maarek. Do you want to take notes?: identifying research missions in yahoo! search pad. In *WWW*, pages 321–330. ACM, 2010. (Cited on page 9.)
 - [76] M. Dredze, J. Blitzer, and F. Pereira. “Sorry, I Forgot the Attachment”: Email attachment prediction. In *CEAS*, 2006. (Cited on pages 10, 29, and 33.)
 - [77] M. Dredze, T. Brooks, J. Carroll, J. Magarick, J. Blitzer, and F. Pereira. Intelligent email: reply and attachment prediction. In *IUI*, pages 321–324. ACM, 2008. (Cited on pages 29 and 33.)
-

- [78] H. Duan and C. Zhai. Mining coordinated intent representation for entity search and recommendation. In *CIKM*, pages 333–342. ACM, 2015. (Cited on pages 14, 89, and 90.)
- [79] H. Duan, C. Zhai, J. Cheng, and A. Gattani. A probabilistic mixture model for mining and analyzing product search log. In *CIKM*, pages 2179–2188. ACM, 2013. (Cited on page 14.)
- [80] H. Duan, C. Zhai, J. Cheng, and A. Gattani. Supporting keyword search in product database: A probabilistic approach. *Proceedings of the VLDB Endowment*, 6(14):1786–1797, 2013. (Cited on page 14.)
- [81] S. T. Dumais. Latent semantic indexing (lsi): Trec-3 report. In *TREC*, pages 219–230. NIST, 1995. (Cited on page 107.)
- [82] H. Fang and C. Zhai. Probabilistic models for expert finding. In *ECIR*, pages 418–430, Berlin, Heidelberg, 2007. Springer-Verlag. (Cited on pages 14 and 51.)
- [83] Y. Fang and A. Godavarthy. Modeling the dynamics of personal expertise. In *SIGIR*, pages 1107–1110, 2014. (Cited on page 69.)
- [84] Y. Fang, L. Si, and A. P. Mathur. Discriminative models of integrating document evidence and document-candidate associations for expert search. In *SIGIR*, pages 683–690, 2010. (Cited on pages 14, 51, 65, and 102.)
- [85] K. Fatahalian, J. Sugerman, and P. Hanrahan. Understanding the efficiency of gpu algorithms for matrix-matrix multiplication. In *SIGGRAPH HWWS*, pages 133–137. ACM, 2004. (Cited on page 68.)
- [86] M. Folk, G. Heber, Q. Kozioł, E. Pourmal, and D. Robinson. An overview of the hdf5 technology suite and its applications. In *EDBT/ICDT Workshop on Array Databases*, pages 36–47. ACM, 2011. (Cited on page 114.)
- [87] I. Forrester Research. US online retail forecast, 2010 to 2015, 2012. (Cited on page 87.)
- [88] M. Gäde, M. Hall, H. Huurdeman, J. Kamps, M. Koolen, M. Skov, E. Toms, and D. Walsh. Overview of the SBS 2015 interactive track. In *CLEF 2015*. Springer, 2015. (Cited on page 87.)
- [89] D. Ganguly, D. Roy, M. Mitra, and G. J. Jones. Word embedding based generalized language model for information retrieval. In *SIGIR*, pages 795–798. ACM, 2015. (Cited on page 13.)
- [90] V. Garcia, E. Debreuve, and M. Barlaud. Fast k nearest neighbor search using gpu. In *CVPRW*, pages 1–6. IEEE, 2008. (Cited on page 107.)
- [91] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, pages 249–256, 2010. (Cited on pages 43, 58, and 96.)
- [92] D. Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Computing Surveys*, 23(1):5–48, 1991. (Cited on page 143.)
- [93] M. Golestan Far, S. Sanner, M. R. Bouadjenek, G. Ferraro, and D. Hawking. On term selection techniques for patent prior art search. In *SIGIR*, pages 803–806. ACM, 2015. (Cited on page 11.)
- [94] D. Graus, D. van Dijk, M. Tsagkias, W. Weerkamp, and M. de Rijke. Recipient recommendation in enterprises using communication graphs and email content. In *SIGIR*, pages 1079–1082. ACM, 2014. (Cited on page 11.)
- [95] D. Graus, M. Tsagkias, W. Weerkamp, E. Meij, and M. de Rijke. Dynamic collective entity representations for entity ranking. In *WSDM*, pages 595–604. ACM, 2016. (Cited on pages 84 and 88.)
- [96] A. Graves and N. Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*, pages 1764–1772, 2014. (Cited on pages 15, 16, and 107.)
- [97] C. Grevet, D. Choi, D. Kumar, and E. Gilbert. Overload is overloaded: email in the age of gmail. In *SIGCHI*, pages 793–802. ACM, 2014. (Cited on page 10.)
- [98] D. A. Grossman, D. O. Holmes, and O. Frieder. A parallel dbms approach to ir. In *TREC*, 1994. (Cited on page 2.)
- [99] D. Guan, H. Yang, and N. Goharian. Effective structured query formulation for session search. Technical report, 2012. (Cited on pages 9 and 20.)
- [100] D. Guan, S. Zhang, and H. Yang. Utilizing query change for session search. In *SIGIR*, pages 453–462. ACM, 2013. (Cited on page 144.)
- [101] C. Gulcehre, M. Moczulski, M. Denil, and Y. Bengio. Noisy activation functions. *arXiv preprint arXiv:1603.00391*, 2016. (Cited on page 112.)
- [102] J. Guo, Y. Fan, Q. Ai, and W. B. Croft. A deep relevance matching model for ad-hoc retrieval. In *CIKM*, pages 55–64. ACM, 2016. (Cited on page 16.)
- [103] J. Guo, Y. Fan, Q. Ai, and W. B. Croft. Semantic matching by non-linear word transportation for information retrieval. In *CIKM*, pages 701–710. ACM, 2016. (Cited on page 13.)
- [104] M. Gupta and M. Bendersky. Information retrieval with verbose queries. *Foundations and Trends® in Information Retrieval*, 9(3-4):209–354, 2015. (Cited on page 12.)
- [105] D. Guthrie, B. Allison, W. Liu, L. Guthrie, and Y. Wilks. A closer look at skip-gram modelling. 2006.

-
- (Cited on page 148.)
- [106] M. Gutmann and A. Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, pages 297–304, 2010. (Cited on pages 92, 93, and 112.)
 - [107] D. Harman. The DARPA TIPSTER project. *SIGIR Forum*, 26(2):26–28, 1992. (Cited on pages 2, 108, and 115.)
 - [108] D. Harman. Document detection data preparation. In *TIPSTER TEXT PROGRAM: PHASE I: Proceedings of a Workshop held at Fredricksburg, Virginia, September 19-23, 1993*, pages 17–31. ACL, 1993. (Cited on pages 115 and 116.)
 - [109] D. Harman and E. Voorhees. Overview of the fifth text retrieval conference. In *TREC-5*, pages 500–238, 1996. (Cited on page 116.)
 - [110] P. E. Hart and J. Graham. Query-free information retrieval. *IEEE Expert*, 12(5):32–37, 1997. (Cited on page 10.)
 - [111] A. Hassan, R. W. White, S. T. Dumais, and Y.-M. Wang. Struggling or exploring?: disambiguating long search sessions. In *WSDM*, pages 53–62. ACM, 2014. (Cited on page 9.)
 - [112] D. Hawking. Overview of the trec-9 web track. In *TREC. NIST*, 2000. (Cited on page 125.)
 - [113] B. He and I. Ounis. Inferring query performance using pre-retrieval predictors. In *SPIRE*, pages 43–54. Springer, 2004. (Cited on pages 12, 37, and 39.)
 - [114] Z. He, S. Liu, M. Li, M. Zhou, L. Zhang, and H. Wang. Learning entity representation for entity disambiguation. In *ACL*, pages 30–34, 2013. (Cited on page 73.)
 - [115] G. E. Hinton. Learning distributed representations of concepts. In *8th Annual Conference of the Cognitive Science Society*, volume 1, page 12, Amherst, MA, 1986. (Cited on pages 15, 51, 53, 73, 88, and 89.)
 - [116] T. Hofmann. Probabilistic latent semantic indexing. In *SIGIR*, pages 50–57. ACM, 1999. (Cited on page 12.)
 - [117] E. Horvitz. Principles of mixed-initiative user interfaces. In *SIGCHI*, pages 159–166. ACM, 1999. (Cited on page 10.)
 - [118] P.-s. Huang, N. M. A. Urbana, X. He, J. Gao, L. Deng, A. Acero, and L. Heck. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*, pages 2333–2338, 2013. (Cited on pages 16, 52, and 54.)
 - [119] S. Huston and W. B. Croft. Evaluating verbose query processing techniques. In *SIGIR*, pages 291–298. ACM, 2010. (Cited on page 11.)
 - [120] B. Huurnink, K. Hofmann, and M. de Rijke. Simulating searches from transaction logs. *Simulation of Interaction*, page 21, 2010. (Cited on page 34.)
 - [121] B. Huurnink, K. Hofmann, M. De Rijke, and M. Bron. Validating query simulators: An experiment using commercial searches and purchases. In *CLEF*, pages 40–51. Springer, 2010. (Cited on page 34.)
 - [122] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *STOC*, pages 604–613. ACM, 1998. (Cited on pages 56 and 68.)
 - [123] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL <http://arxiv.org/abs/1502.03167>. (Cited on pages 112 and 113.)
 - [124] B. J. Jansen and P. R. Molina. The effectiveness of web search engines for retrieving relevant ecommerce links. *Information Processing & Management*, 42(4):1075–1098, 2006. (Cited on pages 14 and 87.)
 - [125] T. Joachims. Optimizing search engines using clickthrough data. In *SIGKDD*, pages 133–142. ACM, 2002. (Cited on pages 34, 42, 98, 103, 104, and 111.)
 - [126] R. Jozefowicz, W. Zaremba, and I. Sutskever. An empirical exploration of recurrent network architectures. In *ICML*, pages 2342–2350, 2015. (Cited on page 15.)
 - [127] A. Kannan, K. Kurach, S. Ravi, T. Kaufmann, A. Tomkins, B. Miklos, G. Corrado, L. Lukács, M. Ganea, P. Young, et al. Smart reply: Automated response suggestion for email. In *KDD*, 2016. (Cited on page 11.)
 - [128] T. Kenter and M. de Rijke. Short text similarity with word embeddings. In *CIKM*, pages 1411–1420. ACM, 2015. (Cited on pages 12 and 77.)
 - [129] T. Kenter, A. Borisov, C. Van Gysel, M. Dehghani, M. de Rijke, and B. Mitra. Neural networks for information retrieval. In *SIGIR 2017*, pages 1403–1406. ACM, 2017.
 - [130] A. M. Kibriya and E. Frank. An empirical comparison of exact nearest neighbour algorithms. In *ECMLPKDD*, pages 140–151. Springer, 2007. (Cited on page 149.)
 - [131] J. Kim and W. B. Croft. Retrieval experiments using pseudo-desktop collections. In *CIKM*, pages 1297–1306. ACM, 2009. (Cited on page 34.)
-

- [132] Y. Kim, J. Seo, and W. B. Croft. Automatic boolean query suggestion for professional search. In *SIGIR*, pages 825–834. ACM, 2011. (Cited on page 11.)
- [133] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014. (Cited on pages 43, 96, and 113.)
- [134] R. Kiros, R. Salakhutdinov, and R. Zemel. Multimodal neural language models. In *ICML*, pages 595–603, 2014. (Cited on page 15.)
- [135] H. Koepke. Why python rocks for research. https://www.stat.washington.edu/~hoytak/_static/papers/why-python.pdf, 2010. Accessed October 13, 2016. (Cited on page 143.)
- [136] A. Kotov, P. N. Bennett, R. W. White, S. T. Dumais, and J. Teevan. Modeling and analysis of cross-session search tasks. In *SIGIR*, pages 5–14. ACM, 2011. (Cited on page 9.)
- [137] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012. (Cited on pages 13, 16, and 108.)
- [138] J. Kruger and D. Dunning. Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *J. Personality and Social Psych.*, 77(6):1121, 1999. (Cited on page 51.)
- [139] J. Krüger and R. Westermann. Linear algebra operators for gpu implementation of numerical algorithms. *ACM Transactions on Graphics*, 22(3):908–916, 2003. (Cited on page 68.)
- [140] G. Kumaran and V. R. Carvalho. Reducing long queries using query quality predictors. In *SIGIR*, pages 564–571. ACM, 2009. (Cited on pages 11, 32, 35, 37, and 39.)
- [141] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *ICML*, pages 1188–1196, 2014. (Cited on pages 13, 73, 74, 75, 91, 108, 110, 114, and 118.)
- [142] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *IEEE*, 86(11):2278–2324, 1998. (Cited on page 13.)
- [143] C.-J. Lee, R.-C. Chen, S.-H. Kao, and P.-J. Cheng. A term dependency-based approach for query terms ranking. In *CIKM*, pages 1267–1276. ACM, 2009. (Cited on pages 11 and 42.)
- [144] O. Levy, Y. Goldberg, and I. Ramat-Gan. Linguistic regularities in sparse and explicit word representations. In *CoNLL*, pages 171–180, 2014. (Cited on page 16.)
- [145] O. Levy, Y. Goldberg, and I. Dagan. Improving distributional similarity with lessons learned from word embeddings. *TACL*, 3:211–225, 2015. (Cited on page 16.)
- [146] H. Li and J. Xu. Semantic matching in search. *Found. & Tr. in Information Retrieval*, 7(5):343–469, June 2014. (Cited on pages 1, 10, 12, 19, 25, 51, 52, 87, 97, 100, and 107.)
- [147] H. Li, W. Liu, and H. Ji. Two-stage hashing for fast document retrieval. In *ACL*. ACL, 2014. (Cited on page 2.)
- [148] S. Liang and M. de Rijke. Formal language models for finding groups of experts. *Information Processing & Management*, 2016. (Cited on page 87.)
- [149] D. J. Liebling, P. N. Bennett, and R. W. White. Anticipatory search: using context to initiate search. In *SIGIR*, pages 1035–1036. ACM, 2012. (Cited on pages 10 and 29.)
- [150] A. Lipani, G. Zuccon, M. Lupu, B. Koopman, and A. Hanbury. The impact of fixed-cost pooling strategies on test collection bias. In *ICTIR*, pages 105–108. ACM, 2016. (Cited on page 125.)
- [151] T.-Y. Liu. *Learning to Rank for Information Retrieval*. Springer, 2011. (Cited on pages 1, 14, 16, 100, and 120.)
- [152] X. Liu, W. B. Croft, and M. Koll. Finding experts in community-based question-answering services. In *CIKM*, pages 315–316. ACM, 2005. (Cited on page 98.)
- [153] E. Loper and S. Bird. NLTK: The natural language toolkit. In *ACL Workshop on Effective Tools and Methodologies for teaching NLP and CL*, pages 63–70. Association for Computational Linguistics, 2002. (Cited on page 143.)
- [154] H. P. Luhn. The automatic creation of literature abstracts. *IBM Journal of R&D*, 2:159–165, 1958. (Cited on pages 75, 76, and 125.)
- [155] J. Luo, X. Dong, and H. Yang. Modeling rich interactions in session search - georgetown university at trec 2014 session track. Technical report, 2014. (Cited on page 9.)
- [156] J. Luo, S. Zhang, and H. Yang. Win-win search: Dual-agent stochastic game in session search. In *SIGIR*, pages 587–596. ACM, 2014. (Cited on page 9.)
- [157] J. Luo, X. Dong, and H. Yang. Session search by direct policy learning. In *ICTIR*, pages 261–270. ACM, 2015. (Cited on pages 9, 22, and 25.)
- [158] M. Lupu, A. Hanbury, et al. Patent retrieval. *Foundations and Trends® in Information Retrieval*, 7(1): 1–97, 2013. (Cited on page 11.)
- [159] C. MacDonald and I. Ounis. Voting for candidates: adapting data fusion techniques for an expert

-
- search task. In *CIKM*, pages 387–396, 2006. (Cited on page 14.)
- [160] C. Macdonald and I. Ounis. Expert search evaluation by supporting documents. In *ECIR*, pages 555–563. Springer, 2008. (Cited on page 13.)
- [161] C. Macdonald and I. Ounis. Usefulness of quality click-through data for training. In *WSCD*, pages 75–79. ACM, 2009. (Cited on page 34.)
- [162] J. B. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, 1967. (Cited on page 78.)
- [163] P. P. Maglio, R. Barrett, C. S. Campbell, and T. Selker. Suitor: An attentive information system. In *IUI*, pages 169–176. ACM, 2000. (Cited on page 10.)
- [164] P. Mahdabi, M. Keikha, S. Gerani, M. Landoni, and F. Crestani. Building queries for prior-art search. In *Information Retrieval Facility Conference*, pages 3–15. Springer, 2011. (Cited on pages 11, 37, 39, and 42.)
- [165] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press. (Cited on pages 2, 9, and 82.)
- [166] I. Matveeva, C. Burges, T. Burkard, A. Laucius, and L. Wong. High accuracy retrieval with multiple nested ranker. In *SIGIR*, pages 437–444. ACM, 2006. (Cited on page 1.)
- [167] K. T. Maxwell and W. B. Croft. Compact query term selection using topically related text. In *SIGIR*, pages 583–592. ACM, 2013. (Cited on page 11.)
- [168] M. T. Maybury. Expert finding systems. Technical Report MTR-06B000040, MITRE, 2006. (Cited on page 13.)
- [169] J. McAuley, R. Pandey, and J. Leskovec. Inferring networks of substitutable and complementary products. In *KDD*, pages 785–794. ACM, 2015. (Cited on page 94.)
- [170] J. McAuley, C. Targett, Q. Shi, and A. van den Hengel. Image-based recommendations on styles and substitutes. In *SIGIR*, pages 43–52. ACM, 2015. (Cited on pages 3 and 94.)
- [171] D. W. McDonald and M. S. Ackerman. Expertise recommender. In *CSCW*, pages 231–240, 2000. (Cited on page 13.)
- [172] S. McPartlin, L. F. Dugal, M. Jenson, and I. W. Kahn. Understanding how US online shoppers are reshaping the retail experience. PricewaterhouseCoopers, 2012. (Cited on page 87.)
- [173] E. Meij, M. Bron, L. Hollink, B. Huurnink, and M. de Rijke. Learning semantic query suggestions. *The Semantic Web-ISWC*, pages 424–440, 2009. (Cited on page 11.)
- [174] D. Metzler and W. B. Croft. Combining the language model and inference network approaches to retrieval. *IPM*, 40(5):735–750, 2004. (Cited on page 22.)
- [175] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *Interspeech*, pages 1045–1048, 2010. (Cited on page 15.)
- [176] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013. (Cited on pages 16, 58, 77, 92, 96, and 112.)
- [177] T. Mikolov, G. Corrado, K. Chen, and J. Dean. Efficient estimation of word representations in vector space. arXiv 1301.3781, 2013. (Cited on pages 15, 16, 53, 73, 88, 91, 92, 97, 107, 108, 110, 114, 118, and 144.)
- [178] T. Mikolov, W.-t. Yih, and G. Zweig. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, pages 746–751, 2013. (Cited on page 16.)
- [179] B. Mitra, F. Diaz, and N. Craswell. Learning to match using local and distributed representations of text for web search. In *WWW*, 2017. (Cited on page 16.)
- [180] M. Mitra, A. Singhal, and C. Buckley. Improving automatic query expansion. In *SIGIR*, pages 206–214. ACM, 1998. (Cited on pages 2 and 36.)
- [181] A. Mnih and G. Hinton. Three new graphical models for statistical language modelling. In *ICML*, pages 641–648, 2007. (Cited on page 15.)
- [182] A. Mnih and G. Hinton. A scalable hierarchical distributed language model. In *NIPS*, pages 1081–1088, 2008. (Cited on page 15.)
- [183] A. Mnih and K. Kavukcuoglu. Learning word embeddings efficiently with noise-contrastive estimation. In *NIPS*, pages 2265–2273, 2013. (Cited on pages 15, 16, 92, and 96.)
- [184] A. Mnih and Y. W. Teh. A fast and simple algorithm for training neural probabilistic language models. In *ICML*, pages 1751–1758, 2012. (Cited on page 92.)
- [185] G. Montavon, G. B. Orr, and K.-R. Müller. *Neural Networks: Tricks of the Trade*. Springer, 2012. (Cited on page 54.)
- [186] G. E. Moore. Cramming more components onto integrated circuits. *Proceedings of the IEEE*, 86(1): 82–85, 1998. (Cited on page 114.)
-

B. Bibliography

- [187] C. Moreira, B. Martins, and P. Calado. Using rank aggregation for expert search in academic digital libraries. In *Simpósio de Informática, INForum*, pages 1–10, 2011. (Cited on pages 14 and 51.)
- [188] M. Muja and D. G. Lowe. Scalable nearest neighbor algorithms for high dimensional data. *Pattern Analysis and Machine Intelligence*, 36(11):2227–2240, 2014. (Cited on page 107.)
- [189] E. D. Mynatt, M. Back, R. Want, M. Baer, and J. B. Ellis. Designing audio aura. In *SIGCHI*, pages 566–573. ACM, 1998. (Cited on page 10.)
- [190] E. Nalisnick, B. Mitra, N. Craswell, and R. Caruana. Improving document ranking with dual word embeddings. In *WWW*, pages 83–84. International World Wide Web Conferences Steering Committee, 2016. (Cited on page 116.)
- [191] A. Y. Ng and M. I. Jordan. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. In *NIPS*, pages 841–848, 2002. (Cited on page 88.)
- [192] R. Nogueira and K. Cho. Task-Oriented Query Reformulation with Reinforcement Learning. *arXiv preprint arXiv:1704.04572*, Apr. 2017. (Cited on page 12.)
- [193] P. Nurmi, E. Lagerspetz, W. Buntine, P. Floréen, and J. Kukkonen. Product retrieval for grocery stores. In *SIGIR*, pages 781–782. ACM, 2008. (Cited on pages 14 and 90.)
- [194] D. Oard, W. Webber, D. Kirsch, and S. Golitsynskiy. Avocado research email collection. *Linguistic Data Consortium*, 2015. (Cited on pages 29, 32, and 41.)
- [195] D. Odijk, E. Meij, I. Sijaranamual, and M. de Rijke. Dynamic query modeling for related content finding. In *SIGIR*, pages 33–42. ACM, 2015. (Cited on page 10.)
- [196] K. D. Onal, I. S. Altıngövdü, P. Karagoz, and M. de Rijke. Getting started with neural models for semantic matching in web search. *arXiv 1611.03305*, 2016. (Cited on page 16.)
- [197] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: bringing order to the web. Technical report, Stanford InfoLab, 1999. (Cited on page 84.)
- [198] M. J. Pazzani and D. Billsus. Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer, 2007. (Cited on page 82.)
- [199] J. Pennington, R. Socher, and C. D. Manning. GloVe: Global Vectors for Word Representation. In *EMNLP*, pages 1532–1543, 2014. (Cited on pages 15, 16, 73, 88, and 110.)
- [200] D. Petkova and W. B. Croft. Hierarchical language models for expert finding in enterprise corpora. In *ICTAI '06*, pages 599–606, 2006. (Cited on pages 14 and 51.)
- [201] S. Petrović, M. Osborne, and V. Lavrenko. Streaming first story detection with application to twitter. In *HLT*, pages 181–189. ACL, 2010. (Cited on page 2.)
- [202] J. Pound, P. Mika, and H. Zaragoza. Ad-hoc object retrieval in the web of data. In *WWW*, pages 771–780. ACM, 2010. (Cited on page 13.)
- [203] W. W. Powell and K. Snellman. The knowledge economy. *Annual review of sociology*, pages 199–220, 2004. (Cited on page 51.)
- [204] L. Prechelt. An empirical comparison of seven programming languages. *Computer*, 33(10):23–29, Oct. 2000. (Cited on page 143.)
- [205] K. Purcell and L. Rainie. Technology’s impact on workers. Technical report, Pew Research Center, 2014. (Cited on page 29.)
- [206] A. Qadir, M. Gamon, P. Pantel, and A. H. Awadallah. Activity modeling in email. In *NAACL-HLT*, pages 1452–1462, 2016. (Cited on page 11.)
- [207] K. Raman, P. N. Bennett, and K. Collins-Thompson. Toward whole-session relevance: exploring intrinsic diversity in web search. In *SIGIR*, pages 463–472. ACM, 2013. (Cited on page 9.)
- [208] R. Řehůřek and P. Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>. (Cited on page 129.)
- [209] R. Reinanda. *Entity Facets for Search*. PhD thesis, Informatics Institute, University of Amsterdam, May 2017. (Cited on page 139.)
- [210] B. Rhodes and T. Starner. Remembrance agent: A continuously running automated information retrieval system. In *PAAMS*, pages 487–495, 1996. (Cited on page 10.)
- [211] B. J. Rhodes. The wearable remembrance agent: A system for augmented memory. In *ISWC*, pages 123–128. IEEE, 1997. (Cited on page 10.)
- [212] B. J. Rhodes. Margin notes: Building a contextually aware associative memory. In *IUI*, pages 219–224. ACM, 2000. (Cited on page 10.)
- [213] B. J. Rhodes and P. Maes. Just-in-time information retrieval agents. *IBM Systems Journal*, 39(3.4): 685–704, 2000. (Cited on page 10.)
- [214] H. M. Robert, S. C. Robert, and D. H. Honemann. *Robert’s rules of order newly revised*. Da Capo Press, 2011. (Cited on page 78.)

-
- [215] S. Robertson. Understanding inverse document frequency: on theoretical arguments for idf. *Journal of documentation*, 60(5):503–520, 2004. (Cited on pages 107 and 115.)
 - [216] S. E. Robertson and S. Walker. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR*, pages 232–241, 1994. (Cited on page 107.)
 - [217] J. Rowley. Product search in e-shopping: a review and research propositions. *Journal of Consumer Marketing*, 17(1):20–35, 2000. (Cited on pages 3, 87, 90, and 97.)
 - [218] D. Rumelhart, G. Hinton, and R. Williams. Learning internal representations by back propagation. In *Parallel Distributed Processing*, pages 318–362. MIT Press, 1986. (Cited on page 54.)
 - [219] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Technical report, DTIC Document, 1985. (Cited on page 15.)
 - [220] N. Ryan, J. Pascoe, and D. Morse. Enhanced reality fieldwork: the context aware archaeological assistant. In *CAA*, pages 269–274. Archaeopress, 1999. (Cited on page 10.)
 - [221] J. Rybak, K. Balog, and K. Nørkvåg. Temporal expertise profiling. In *ECIR*, pages 540–546. Springer, 2014. (Cited on page 69.)
 - [222] H. Sak, A. W. Senior, and F. Beaufays. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Interspeech*, 2014. (Cited on page 15.)
 - [223] R. Salakhutdinov and G. Hinton. Semantic hashing. *Int. J. Approximate Reasoning*, 50(7):969–978, 2009. (Cited on pages 13, 51, and 54.)
 - [224] G. Salton, C. Buckley, and E. A. Fox. Automatic query formulations in information retrieval. *JASIST*, 34(4):262, 1983. (Cited on pages 35 and 36.)
 - [225] N. Sawhney and C. Schmandt. Nomadic radio: speech and audio interaction for contextual messaging in nomadic environments. *TOCHI*, 7(3):353–383, 2000. (Cited on page 10.)
 - [226] A. Schuth. Search engines that learn from their users. *SIGIR Forum*, 50(1):95–96, 2016. (Cited on page 84.)
 - [227] D. Sculley and G. Inc. Large scale learning to rank. In *In NIPS 2009 Workshop on Advances in Ranking*, 2009. (Cited on pages 42, 98, and 103.)
 - [228] P. Serdyukov and D. Hiemstra. Modeling documents as mixtures of persons for expert finding. In *ECIR*, pages 309–320. Springer, 2008. (Cited on page 14.)
 - [229] P. Serdyukov, H. Rode, and D. Hiemstra. Modeling multi-step relevance propagation for expert finding. In *CIKM*, pages 1133–1142, 2008. (Cited on page 14.)
 - [230] C. Shannon. A mathematical theory of communication. *Bell System Technical J.*, 27:379–423, 623–656, 1948. (Cited on page 61.)
 - [231] J. A. Shaw, E. A. Fox, J. A. Shaw, and E. A. Fox. Combination of multiple searches. In *TREC*, pages 243–252, 1994. (Cited on pages 65 and 120.)
 - [232] Y. Shen, X. He, J. Gao, L. Deng, and G. Mesnil. A latent semantic model with convolutional-pooling structure for information retrieval. In *CIKM*, pages 101–110, 2014. (Cited on page 16.)
 - [233] M. Shokouhi and Q. Guo. From queries to cards: Re-ranking proactive card recommendations based on reactive search history. In *SIGIR*, pages 695–704. ACM, 2015. (Cited on pages 10 and 29.)
 - [234] M. D. Smucker, J. Allan, and B. Carterette. A comparison of statistical significance tests for information retrieval evaluation. In *CIKM*, pages 623–632. ACM, 2007. (Cited on pages 44, 58, 59, 67, 97, and 119.)
 - [235] Y. Song and Q. Guo. Query-less: Predicting task repetition for nextgen proactive search and recommendation engines. In *WWW*, pages 543–553. International World Wide Web Conferences Steering Committee, 2016. (Cited on pages 10 and 29.)
 - [236] P. Sorg and P. Cimiano. Finding the right expert: Discriminative models for expert retrieval. In *KDIR*, pages 190–199, 2011. (Cited on page 14.)
 - [237] K. Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21, 1972. (Cited on page 115.)
 - [238] T. Strohman, D. Metzler, H. Turtle, and W. B. Croft. Indri: A language model-based search engine for complex queries. In *ICIA*, 2005. (Cited on pages 7, 42, 116, and 143.)
 - [239] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, pages 3104–3112, 2014. (Cited on pages 13 and 15.)
 - [240] J. Tague, M. Nelson, and H. Wu. Problems in the simulation of bibliographic retrieval systems. In *SIGIR*, pages 236–255. Butterworth & Co., 1980. (Cited on page 34.)
 - [241] J. M. Tague and M. J. Nelson. Simulation of user judgments in bibliographic retrieval systems. In *ACM SIGIR Forum*, volume 16, pages 66–71. ACM, 1981. (Cited on page 34.)
 - [242] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016. URL <http://arxiv.org/abs/1605.02688>. (Cited on pages 147, 148, and 149.)
-

B. Bibliography

- [243] K. Tran, A. Bisazza, and C. Monz. Recurrent memory network for language modeling. In *NAACL*, pages 321–331, 2016. (Cited on page 15.)
- [244] TREC. TREC1-8 Adhoc Track, 1992–1999. (Cited on pages 108, 115, and 125.)
- [245] TREC. Enterprise Track, 2005–2008. (Cited on pages 3, 13, 51, and 56.)
- [246] TREC. Session Track, 2009–2014. (Cited on pages 9, 22, and 25.)
- [247] X. Tu, J. X. Huang, J. Luo, and T. He. Exploiting semantic coherence features for information retrieval. In *SIGIR*, pages 837–840. ACM, 2016. (Cited on page 13.)
- [248] J. Turian, L. Ratinov, and Y. Bengio. Word representations: a simple and general method for semi-supervised learning. In *ACL*, pages 384–394, 2010. (Cited on pages 15, 73, and 88.)
- [249] A. K. Uysal and S. Gunal. The impact of preprocessing on text classification. *Information Processing & Management*, 50(1):104–112, 2014. (Cited on page 143.)
- [250] S. van der Walt, S. C. Colbert, and G. Varoquaux. The numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011. (Cited on page 148.)
- [251] D. van Dijk, M. Tsagkias, and M. de Rijke. Early detection of topical expertise in community question and answering. In *SIGIR*, 2015. (Cited on page 51.)
- [252] C. Van Gysel. Listening to the flock - towards opinion mining through data-parallel, semi-supervised learning on social graphs. Master’s thesis, University of Antwerp, 2014.
- [253] C. Van Gysel, M. de Rijke, and M. Worring. Semantic entities. In *ESAIR*, pages 1–2. ACM, 2015. (Cited on pages 6 and 7.)
- [254] C. Van Gysel, B. Goethals, and M. de Rijke. Determining the presence of political parties in social circles. In *ICWSM*, volume 2015, pages 690–693, 2015.
- [255] C. Van Gysel, L. Velikovich, I. McGraw, and F. Beaufays. Garbage modeling for on-device speech recognition. In *Interspeech*, volume 2015, pages 2127–2131, 2015.
- [256] C. Van Gysel, M. de Rijke, and E. Kanoulas. Learning latent vector spaces for product search. In *CIKM*, pages 165–174. ACM, 2016. (Cited on pages 6, 7, 108, 109, 110, 112, 114, 118, 119, 126, 128, 147, 148, 149, and 150.)
- [257] C. Van Gysel, M. de Rijke, and M. Worring. Unsupervised, efficient and semantic expertise retrieval. In *WWW*, pages 1069–1079. ACM, 2016. (Cited on pages 7, 76, 77, 78, 84, 87, 88, 92, 100, 110, 126, 128, 147, 148, and 150.)
- [258] C. Van Gysel, E. Kanoulas, and M. de Rijke. Lexical query modeling in session search. In *ICTIR*, pages 69–72. ACM, 2016. (Cited on pages 6 and 144.)
- [259] C. Van Gysel, M. de Rijke, and E. Kanoulas. Structural regularities in expert vector spaces. In *ICTIR*. ACM, 2017. (Cited on pages 7, 147, and 150.)
- [260] C. Van Gysel, M. de Rijke, and E. Kanoulas. Neural vector spaces for unsupervised information retrieval. Under review, 2017. (Cited on page 7.)
- [261] C. Van Gysel, M. de Rijke, and E. Kanoulas. Semantic entity retrieval toolkit. In *Neu-IR SIGIR Workshop*, 2017. (Cited on pages 8 and 77.)
- [262] C. Van Gysel, E. Kanoulas, and M. de Rijke. Pyndri: a python interface to the indri search engine. In *ECIR*, volume 2017. Springer, 2017. (Cited on pages 7, 42, 116, and 150.)
- [263] C. Van Gysel, B. Mitra, M. Venanzi, R. Rosemarin, G. Kukla, P. Grudzien, and N. Cancedda. Reply with: Proactive recommendation of email attachments. In *CIKM*, 2017. (Cited on pages 6 and 7.)
- [264] C. Van Gysel, I. Oparin, X. Niu, and Y. Su. Rank-reduced token representation for automatic speech recognition, 2017. US Patent Application 15/459,481.
- [265] C. J. van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann, 2nd edition, 1979. (Cited on page 112.)
- [266] V. Vapnik. *Statistical learning theory*, volume 1. Wiley New York, 1998. (Cited on pages 51, 88, and 98.)
- [267] M. Vidal-Naquet and S. Ullman. Object recognition with informative features and linear classification. In *ICCV*, page 281. IEEE, 2003. (Cited on page 108.)
- [268] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *JMLR*, 11:2837–2854, 2010. (Cited on page 80.)
- [269] E. M. Voorhees. The TREC robust retrieval track. *SIGIR Forum*, 39(1):11–20, June 2005. (Cited on pages 115 and 116.)
- [270] I. Vulić and M.-F. Moens. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In *SIGIR*, pages 363–372. ACM, 2015. (Cited on pages 12, 108, 110, 114, and 118.)
- [271] R. Weber, H.-J. Schek, and S. Blott. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In *VLDB*, pages 194–205, 1998. (Cited on pages 75, 99,

-
- and 105.)
- [272] W. Weerkamp, K. Balog, and M. de Rijke. Using contextual information to improve search in email archives. In *ECIR*, pages 400–411. Springer, 2009. (Cited on page 42.)
 - [273] X. Wei and W. B. Croft. Lda-based document models for ad-hoc retrieval. In *SIGIR*, pages 178–185. ACM, 2006. (Cited on page 109.)
 - [274] S. A. Weil, D. Tinapple, and D. D. Woods. New approaches to overcoming e-mail overload. In *HFES*, volume 48, pages 547–551. SAGE, 2004. (Cited on pages 35 and 44.)
 - [275] S. Whittaker and C. Sidner. Email overload: Exploring personal information management of email. In *SIGCHI*, pages 276–283. ACM, 1996. (Cited on page 10.)
 - [276] Wikipedia. List of nvidia graphics processing units — wikipedia, the free encyclopedia, 2017. URL https://en.wikipedia.org/w/index.php?title=List_of_Nvidia_graphics_processing_units&oldid=792964538. [Online; accessed 8-August-2017]. (Cited on page 114.)
 - [277] J. Xu, C. Chen, G. Xu, H. Li, and E. R. T. Abib. Improving quality of training data for learning to rank using click-through data. In *WSDM*, pages 171–180. ACM, 2010. (Cited on page 34.)
 - [278] X. Xue and W. B. Croft. Automatic query generation for patent search. In *CIKM*, pages 2037–2040. ACM, 2009. (Cited on pages 11, 37, 39, and 42.)
 - [279] X. Xue and W. B. Croft. Transforming patents into prior-art queries. In *SIGIR*, pages 808–809. ACM, 2009. (Cited on pages 11 and 42.)
 - [280] X. Xue, S. Huston, and W. B. Croft. Improving verbose queries using subset distribution. In *CIKM*, pages 1059–1068. ACM, 2010. (Cited on pages 11, 32, and 35.)
 - [281] H. Yang, D. Guan, and S. Zhang. The query change model: Modeling session search as a markov decision process. *TOIS*, 33(4):20:1–20:33, 2015. (Cited on pages 9, 20, 22, and 25.)
 - [282] H. Zamani and W. B. Croft. Embedding-based query language models. In *ICTIR*, pages 147–156. ACM, 2016. (Cited on page 13.)
 - [283] H. Zamani and W. B. Croft. Estimating embedding vectors for queries. In *ICTIR*, pages 123–132. ACM, 2016. (Cited on page 13.)
 - [284] M. D. Zeiler. Adadelta: An adaptive learning rate method. *CoRR*, abs/1212.5701, 2012. (Cited on page 58.)
 - [285] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to ad hoc information retrieval. In *SIGIR*, pages 334–342. ACM, 2001. (Cited on page 144.)
 - [286] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to information retrieval. *TOIS*, 22(2):179–214, 2004. (Cited on pages 42, 48, 98, 100, 107, 108, 109, 110, 115, 119, and 138.)
 - [287] L. Zhao and J. Callan. Term necessity prediction. In *CIKM*, pages 259–268. ACM, 2010. (Cited on page 12.)
 - [288] Y. Zhao, F. Scholer, and Y. Tsegay. Effective pre-retrieval query performance prediction using similarity and variability evidence. In *ECIR*, pages 52–64. Springer, 2008. (Cited on pages 37 and 39.)
 - [289] Y. Zhao, L. Zhiyuan, and M. Sun. Representation learning for measuring entity relatedness with rich information. In *IJCAI*, pages 1412–1418, 2015. (Cited on pages 15, 73, 88, and 139.)
 - [290] G. Zuccon, B. Koopman, P. Bruza, and L. Azzopardi. Integrating and evaluating neural word embeddings in information retrieval. In *20th Australasian Document Computing Symposium*, pages 12:1–12:8. ACM, 2015. (Cited on pages 13 and 116.)
-

Samenvatting

Zoekmachines zijn enorm afhankelijk van methodes gebaseerd op overlappende lexicale termen: methodes die zoekopdrachten en documenten voorstellen als een verzameling van woorden en die de gelijkenis tussen zoekopdracht en een document berekenen aan de hand van zoektermen die exact voorkomen in het document. Tekst—een document of een zoekopdracht—wordt voorgesteld als een zak van de woorden die voorkomen in de tekst, waarbij grammatica of woordvolgorde genegeerd wordt, maar de woordfrequentie behouden blijft. Wanneer de gebruiker een zoekopdracht opgeeft, sorteert de zoekmachine documenten aan de hand van een relevantie score, die onder meer bepaald wordt door de mate van overeenkomst tussen termen die voorkomen in de zoekopdracht en het document. Hoewel methodes gebaseerd op lexicale termen intuïtief en effectief zijn in de praktijk, steunen ze erg hard op de hypothese dat documenten waarin de zoekopdracht exact voorkomt relevant zijn voor de zoekopdracht, ongeacht de betekenis van de zoekopdracht en de bedoelingen van de gebruiker. Omgekeerd, methodes gebaseerd op lexicale termen veronderstellen dat documenten die geen enkele zoekopdracht term bevatten irrelevant zijn tot de zoekopdracht van de gebruiker. Maar het is bekend dat een hoge graad van overeenkomst op het term-niveau niet noodzakelijk relevantie impliceert, en omgekeerd, dat documenten die geen zoekopdracht termen bevatten wel relevant kunnen zijn. Bijgevolg bestaat er een *woordenschat kloof* tussen zoekopdrachten en documenten die voorkomt als beide verschillende termen gebruiken om dezelfde concepten te beschrijven. Het is het bestrijden van het effect dat voortgebracht wordt door deze woordenschat kloof dat het onderwerp is van deze dissertatie.

In het eerste deel van deze dissertatie formuleren we zoekopdrachten—voor het ophalen van documenten met methodes gebaseerd op lexicale term overeenkomst—van complexe en heterogene tekstuele structuren (zoeksessies en email conversaties) om zo goed mogelijk de informatiebehoefte van de gebruiker te vervullen. In dit scenario is het gebruiken van de volledige tekstuele structuur als zoekopdracht (a) computationeel kostbaar gezien de lange lengte van de zoekopdracht en (b) gevoelig tot het onjuist classificeren van document als irrelevant door de aanwezigheid van vervuilende termen in de tekstuele structuur die niet voorkomen in relevant documenten. Zoeksessies bestaan uit een opeenvolging van gebruikersinteracties, resultaten getoond door de zoekmachine en zijn gewoonlijk het teken van een complexe informatiebehoefte van de gebruiker. Het is dan de bedoeling om een tekstuele zoekopdracht te formuleren die de informatiebehoefte kan vervullen en de sessie voldoende beschrijft. In het geval van email conversaties is de informatiebehoefte van de gebruiker minder duidelijk. We richten ons op het specifieke geval waar een binnenkomende email een vraag voor inhoud, bijvoorbeeld een document of een hyperlink, bevat. We willen vervolgens een zoekopdracht formuleren die het correcte item ophaalt uit een verzameling documenten. Het doel van het onderzoek van het eerste deel van deze dissertatie is het formuleren van zoekopdrachten die (a) korter zijn en (b) betere resultaten ophalen dan als we de gehele tekstuele structuur gebruiken als zoekopdracht.

In het tweede deel van deze dissertatie, vermijden we methodes die gebaseerd zijn op de overeenkomst van lexicale termen tussen zoekopdrachten en documenten. Dit staat ons toe om de woordenschat kloof te vermijden door te vergelijken op basis van

semantische concepten in plaats van lexicale termen. Documenten en zoekopdrachten worden voorgesteld als vectoren in een ruimte van lage dimensionaliteit. Een zoekopdracht, die door de gebruiker uitgedrukt wordt in termen, wordt geprojecteerd naar een latente vectorruimte van zoekopdrachten. Vervolgens wordt de latente representatie van de zoekopdracht geprojecteerd naar de latente ruimte van documenten. Documenten worden vervolgens gesorteerd in afnemende volgorde van de gelijkenis van hun latente representatie met de geprojecteerde representatie van de zoekopdracht. We focussen op zoekdomeinen waarvan bekend is dat semantisch vergelijken belangrijk is: entiteiten en nieuwsartikelen.

Summary

Search engines rely heavily on term-based approaches that represent queries and documents as bags of words. Text—a document or a query—is represented by a bag of its words that ignores grammar and word order, but retains word frequency counts. When presented with a search query, the engine then ranks documents according to their relevance scores by computing, among other things, the matching degrees between query and document terms. While term-based approaches are intuitive and effective in practice, they are based on the hypothesis that documents that exactly contain the query terms are highly relevant regardless of query semantics. Inversely, term-based approaches assume documents that do not contain query terms as irrelevant. However, it is known that a high matching degree at the term level does not necessarily mean high relevance and, vice versa, documents that match null query terms may still be relevant. Consequently, there exists a *vocabulary gap* between queries and documents that occurs when both use different words to describe the same concepts. It is the alleviation of the effect brought forward by this vocabulary gap that is the topic of this dissertation.

In the first part of this dissertation, we formulate queries—for retrieval using term-based approaches—from complex and heterogeneous textual structures (search sessions and email threads) so as to fulfill the user’s information need. In this scenario, issuing the full textual structure as a query is (a) computationally expensive due to the long query length and (b) prone to classifying documents as false negatives for retrieval due to noisy terms part of the textual structure that do not occur in relevant documents. Search sessions consist of a sequence of user interactions, search engine responses and are indicative of a complex information need. The task is then to formulate a textual query to satisfy the overall information need that characterizes the session. In the case of email threads, the explicit information need is less straightforward. We focus on a particular case where an incoming email message is a request for content. The task is then to formulate a query that correctly retrieves the appropriate item from a document repository. Overall, the common goal of the research performed in first part of this dissertation is to formulate queries that are (a) shorter and (b) exhibit more effective retrieval than taking the full textual structure as a query.

In the second part of this dissertation, we steer away from term-based approaches. This allows us to avoid the vocabulary gap altogether by matching based on semantic concepts rather than lexical terms. Documents and queries are represented by low-dimensional representations in a latent vector space. Term-based queries—entered by the user—are first projected into the low-dimensional query space. Subsequently, the latent representation of the user query is projected to the latent space of documents. Documents are then ranked in decreasing order of similarity between the document’s representation and the user query representation. We focus on retrieval domains where semantic matching is known to be important: entities and news articles.

Search engines rely heavily on term-based approaches that represent queries and documents as bags of words. Text—a document or a query—is represented by a bag of its words that ignores grammar and word order, but retains word frequency counts. When presented with a search query, the engine then ranks documents according to their relevance scores by computing, among other things, the matching degrees between query and document terms. While term-based approaches are intuitive and effective in practice, they are based on the hypothesis that documents that exactly contain the query terms are highly relevant regardless of query semantics. Inversely, term-based approaches assume documents that do not contain query terms as irrelevant. However, it is known that a high matching degree at the term level does not necessarily mean high relevance and, vice versa, documents that match null query terms may still be relevant. Consequently, there exists a vocabulary gap between queries and documents that occurs when both use different words to describe the same concepts. It is the alleviation of the effect brought forward by this vocabulary gap that is the topic of this dissertation. More specifically, we propose (1) methods to formulate an effective query from complex textual structures and (2) latent vector space models that circumvent the vocabulary gap in information retrieval.

ISBN 9789461828385



9 789461 828385